ROADVIEW

Robust Automated Driving in Extreme Weather

Project No. 101069576

Deliverable D4.5

Initial readiness assessment of specific datasets

WP4 Data logging and readiness

Authors	Ian Marsh (RISE)
Lead participant	RISE, Research Institutes of Sweden, AB
Delivery date	21 May 2024
Dissemination level	Public
Type	Document

Version 04

Co-funded by the European Union

Innovate UK

Co-funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Climate, Infrastructure and Environment Executive Agency (CINEA). Neither the European Union nor the granting authority can be held responsible for them. Project grant no. 101069576.

UK participants in this project are co-funded by Innovate UK under contract no.10045139.
Swiss participants in this project are co-funded by the Swiss State Secretariat for Education, Research and Innovation (SERI) under contract no. 22.00123.

Revision history

Author(s)	Description	Date
Ian Marsh (RISE)	Draft V0.0	24/04/2023
Ian Marsh (RISE)	Draft V0.1	01/06/2023
Ian Marsh (RISE)	Draft V0.2	21/07/2023
Ian Marsh (RISE)	Draft V0.3 (sent out)	20/08/2023
Ian Marsh (RISE)	Draft V0.3a (working version)	24/08/2023
Ian Marsh (RISE)	Draft V0.4	25/08/2023
Ian Marsh (RISE)	Draft V0.5	27/08/2023
Joachim Glaess (KO)	Revision	28/08/2023
Ian Marsh (RISE)	Draft V0.6	29/08/2023
Ian Marsh (RISE)	Draft V0.7	29/08/2023
Ian Marsh (RISE)	Draft V0.8	30/08/2023
Heikki Hyyti (FGI), Eren Aksoy (HH)	Revision	30/08/2023
Eren Aksoy (HH), Ian Marsh (RISE)	Final version (v01)	31/08/2023
Andreia Cruz (accelCH)	Formatting and formal check (v01)	31/08/2023
Ian Marsh (RISE)	Version 02	08/05/2024 Reviewers Comments This document reviews several specific datasets that the project will need for its operation, including data relating to sensors, to road scenes, and weather models. These are loosely tied to the reference architecture presented in D2.3.

		However, there is no clear analysis, based on this architecture and the project methodology, of what specific (logical) datasets are required. The "readiness" of datasets for the project cannot be determined without this. <u>Responses and changes</u> 1. Page 13 summarises responses to the issues above. They are in this location, as fixes are in relation to the functional architecture. 2. Page 19-20, Table 2 shows the minimum setup for ROADVIEW. The sensor setup will have a bearing on the quality, hence its inclusion and helps explain the methodology. 3. Page 21, Table 3 has been expanded to show which tools we are using to derive the DRLs for image and LiDAR quality. 4. Page 23, Figure 7 has been added for clarity. 5. Pages 30-31 explains the use of DRLs in some detail. This exemplifies the readiness of datasets in concrete terms. 6. Page 31, Figure 12 is new to show the route taken in FGIs rural journey. FGI is the Finnish Geospatial Institute's dataset. 7. Page 33, Figure 14 shows the DRL applied to FGI dataset. 8. Page 34, a summary of DRLs and work ahead. 9. Page 50, a new reference.
Ian Marsh (RISE)	Version 03	Missing line on page 23 added.
Mario Ceccarelli (accelCH)	Logo & funding acknowledgement update	21/05/2024

Contents

List of Figures

List of Tables

Partner short names

HH	Hogskölan Halmstad
LUA	Lapin Ammattikorkeakoulu OY
THI	Technische Hochschule Ingolstadt
VTI	Statens Vag och Transportforskningsinstitut
CE	Centre d'études et d'expertise sur les risques, l'environnement, la mobilité et l'aménagement
RISE	RISE Research Institutes of Sweden AB
FGI	Maanmittauslaitos – Finnish Geospatial Research Institute
Repli5	Repli5
KO	Konrad GmbH
FORD	Ford Otomotiv Sanayi Arnonim Sirketi
CRF	Canon Research Centre France
accelCH	accelopment Schweiz AG
WMG	The University of Warwick
VTT	Teknologian Tutkimuskeskus VTT OY

Abbreviations

AD	Automated Driving
CAV	Connected Automated Vehicle
D	Deliverable
DDI	Direct Data Injection
DRL	Data Readiness Level
EC	European Commission
EU	European Union
GA	Grant Agreement
GNSS	Global Navigation Satellite System
HD	High Definition
HD MAP	Detailed and highly accurate digital representation of the real-world environment
HEU	Horizon Europe
HiL	Hardware-in-the-Loop
IMU	Inertial Measurement Unit
M	Month
ML	Machine Learning
MS	Milestone
MRM	Minimum Risk Manoeuvre
ODD	Operational Design Domain
OEM	Original Equipment Manufacturer
OTA	Over-the-Air
TRL	Technology Readiness Level
ViL	Vehicle-in-the-Loop
WP	Work Package

Executive summary

The outcome of this deliverable is a description of the data readiness levels (DRLs) in the ROADVIEW project. Essentially DRLs are the data 'equivalent' of Technical Readiness Levels (TRLs). Specific is included in the title means internally (generated) and externally (available) datasets for autonomous driving. As well as the datasets we introduce the concept of data readiness levels and how the levels are implemented. DRLs are a general concept for differing data kinds, that said, we will use our own collated metrics for the specifics related to autonomous driving initially. Data quality detractors are not solely sensor related, non-domain aspects, missing data, corrupted values, processing errors are generic problems in data processing and can lead to insufficient HD map quality and therein can miss important objects in the test drives. We do not provide a DRL value for the datasets in this deliverable, as LiDAR and RADAR data has not been evaluated in sufficient detail, but we provide a roadmap for LiDAR evaluation.

Clearly, data for autonomous driving is an enormous area crossing many disciplines, vision, ML, metrological, measurement techniques, sensing, instrumentation, and we do not attempt to cover all these topics however provide a slice through the topic with an example of image quality.

Objectives

The main objective of this deliverable is a report on how to quantitatively evaluate the Data Readiness Levels (DRL) of datasets. With the help of this deliverable, we aim to assign a score, 1 (low) to 9 (high) of a specific dataset we have available to us. Datasets can be ROADVIEW partner produced (page 20) or openly available ones (page 24). The final project goal is the ability to upload a dataset and an evaluation thereof of that data for autonomous driving.

Methodology and implementation

The main methodology followed is to define a simple, working assessment of the datasets used in ROADVIEW. Irrespective of the algorithms' performance (F-Scores, Confusion matrices and so on) the success of autonomous vehicles will largely depend on the quality of data used. Therefore, the data should be of sufficient quality for the scenarios and settings defined in the project. Since the goal of the project is to enable autonomous vehicles to operate in rain and snow, the coverage, quality, and settings will be a determining factor in the projects' success. Also, quantitively, where the operation of the vehicle perception is too low, we can pinpoint at least where the data is/was insufficient using the main concept. Backtracing from an object misclassification (or missing at all) to the HD MAP, data source (Visible camera, Infra-red, LiDAR, RADAR) to the pixels and potentially a sensor issue would be a truly valuable tool in this field, akin to a debugger in software development.

Outcome

The outcome of this first deliverable on this topic is an overview of the domain of road sensor fused simple. It is more of an informational type of deliverable, however, with concrete examples from internal and external datasets.

Next steps

The next steps are to continue the development of the DRL concepts and obtain a RADAR set for the same datasets. We will also assess the public datasets in further detail, by adding additional quality metrics. Assessments must be developed or in the public domain. The external datasets in our focus are `nuScenes` and `mmDetection3D`.

Part I: Preliminaries

What's in this deliverable?

Essentially this deliverable contains an evaluation of internally produced and publicly available datasets for autonomous driving. An emphasis is on scenes with poorer driving conditions. How driving conditions impacts on sensor performance and subsequent data flows is also addressed. An end-to-end software example is used as an illustration. VISIBLE and thermal cameras are considered, RADAR and ultrasonic sensors less so.

What's not in this deliverable?

Some issues are not included in this deliverable, on is sensor calibration, cameras and LiDARs (though LiDAR calibration can be found in this Link). This is mostly due to the competence area, as well as access to specialised equipment. We will also not discuss issues in ROADVIEW available in other Work packages and deliverables however are related to datasets. Examples include ODDs (WP2, Task 2.1) and the Data Management Plan but (WP1, T1.4) refer to other work packages and tasks within ROADVIEW. RADAR data is not discussed in detail this deliverable, again due to competence in this task and data at a timely point. It will be covered in future dataset evaluation.

Data in ROADVIEW: a perspective

Introduction

In this first deliverable (of 3) on data quality, we split the initial data readiness into 3 sections:

1. Introduce data readiness concepts.
2. The data details about the data set(s) for assessment.
3. Highlight practicalities around storage, processing, formats, sizes, and costs.

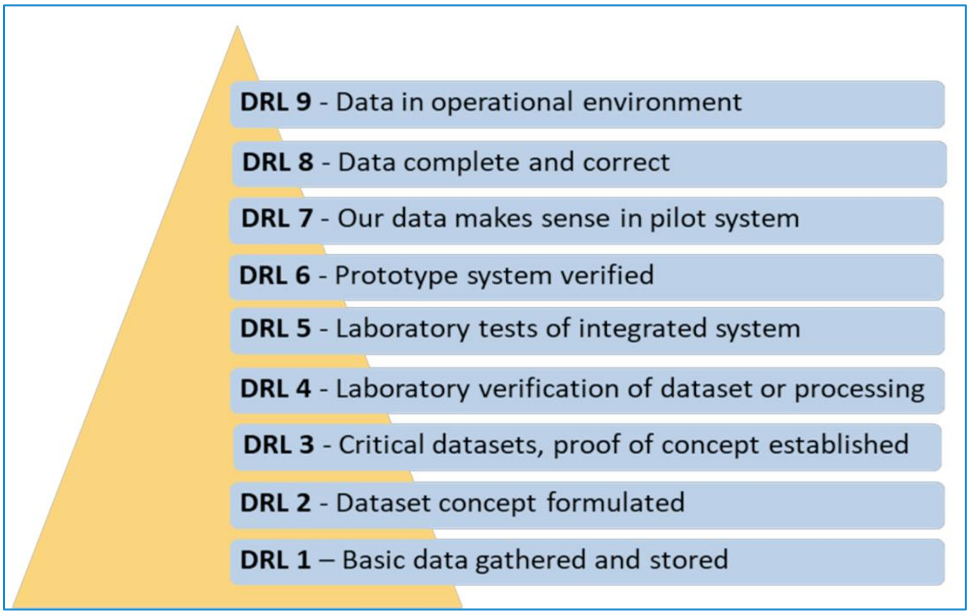

Figure 1 - Data plays a central role in the many projects. Above, is a generic view of data of the data processing stages. The lower DRLs are processed bottom first. A specific DRL illustration is in Fig. 8.

The big (data) picture

We look at the increasing role of data in computer-based tasks. Taking a step back computer science works on some simple principles. Although many computer scientists developed these, the idea of splitting the thinking process into Data and Control was coined by Alfred Aho (still lectures at 82 years old at Columbia University) of AWK fame). These are shown below. Looking back, the amount of data processed was much less than now and the focus was on controlling the program to achieve complex tasks. Data was typically input → processed → output.

Programs = Control + Data

As key algorithms (many formalised and popularised by Knuth) development of the algorithms, especially their complexity became key. Control flow became algorithms, which can be used again and again in different contexts. Again, the separation of the data and the algorithms is a key philosophy. Many systems and programs today could benefit from working code to a specific (well known) algorithm that can be changed or improved. Searching and sorting being the classics.

Application = Algorithms + Data

Jumping forward a few decades in computer science, a key change has been the use of large amounts of data. This jump in huge data amounts is not only related to Machine Learning and AI, but gathering data from our environment, roads, earth observations, Internet logs, nuclear particles and so on. Whence scanning, capturing, measuring our environment (at whatever timescale or physical dimensions) data has become pervasive. In autonomous driving, sensing, gathering, processing, validating, (even legalising) data is so important, it is almost the most important facet of a project such as ROADVIEW. We do not demean the algorithms or computer science; however, the roll of data is of uppermost attention in projects such as ROADVIEW. It is logical that the quality of the data, should be and is an important part of the project.

Autonomous driving = Algorithms + Training Data + Testing Data + Testing

Therefore, ROADVIEW will (i) introduce a novel concept of sensor denoising to filter out noisy sensor readings (camera, LiDAR, and RADAR) and (ii) assign a data readiness level to validate the quality of the raw data before passing them to the perception modules. This will lead to more robust perception under varying environments and weather conditions. In terms of assessment, we need to stipulate that in image and video quality for ML processing (pipelines) introduces new and different issues when quality assessment is for people. This we will come back to when discussing the VISIBLE image and video quality.

The ROADVIEW reference architecture (a summary of D2.3)

The work in this task is in essence adjunct to the architecture work done in the Architecture Reference Task, part of WP2. Figure 2 below shows the generic architecture ROADVIEW uses. It represents the core functional generic architecture of the ROADVIEW project. Novel ROADVIEW modules are shown in dark grey, and available OEM modules in light grey colours, respectively. There are 3 functionalities: Perception, Planning, and Control in red, purple, & blue. The **perception** block functions are for robust environment perception. Basically, filtering of noisy sensor data, low-level sensor fusion, object detection, free-space detection, weather-type detection, slipperiness detection, and visibility detection. The perception block receives input from various sensing modalities such as RGB cameras, LiDARs, RADARs, Thermal cameras, Inertial Measurement Units (IMU), and Global Navigation Satellite System (GNSS) modules. The **planning** block is for localisation, trajectory prediction & planning functionalities. The **controller** block focuses on weather-aware decision-making and control of the velocity and acceleration-related parameters. Each sensor reading is processed individually up to the low-level sensor fusion module, which feeds the downstream perception modules such as object detection. Dashed arrows in the figure represented sensor readings cleaned from noise (i.e., outliers) introduced by, for instance, falling snow particles or rain drops.

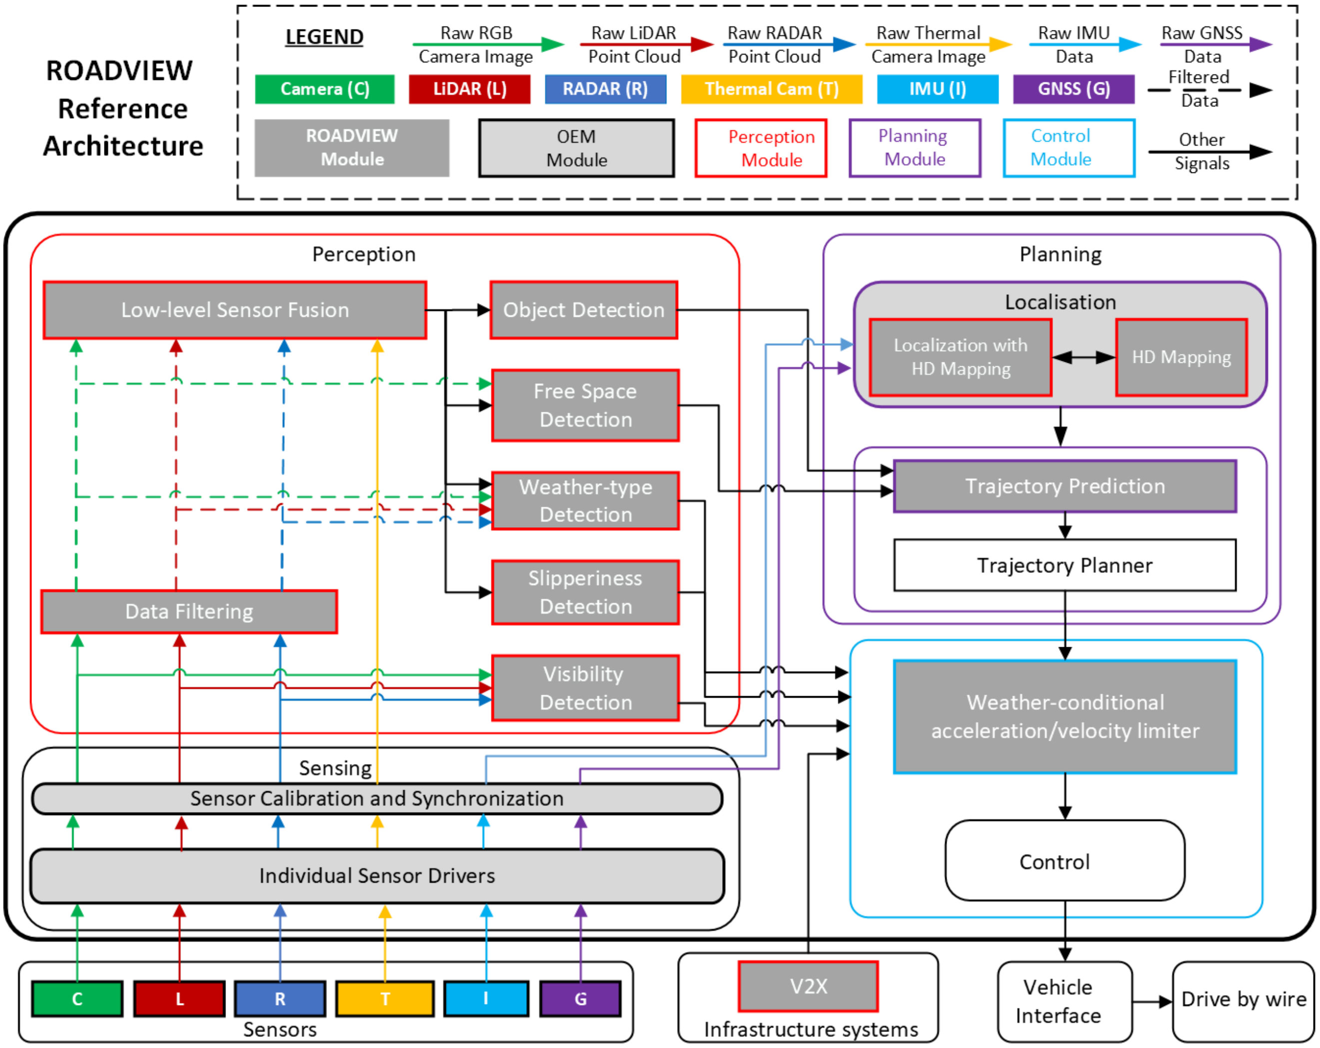

Figure 2 - The ROADVIEW reference architecture (source D2.3)

The **reference architecture** in Figure 2 drives our Readiness Level development. Table 3 shows the DRLs, and the sensor calibration and synchronisation above correspond to level 0 in our DRLs. That is before data is gathered, and per modality, the calibration needs to be done. We mean extrinsic calibration with respect to calibration in this case. Data filtering maps to data quality assessment in our DRL schema, that is levels 2-7. These are described in more detail below with an example. We do not look at the quality of the fused modalities in DRL in this first report, rather per modality.

The **project methodology** with respect to data is to find the best from each dataset and as a final dataset release a ROADVIEW dataset. Indeed, one of the objectives of ROADVIEW is 'stamp' the dataset with a quality value from 1 (lowest) to 9 (highest). We need to state here, the value attributed to a dataset, or indeed amalgamation should not be used to compare datasets at this stage. That is because the setup, calibrations, vehicles, conditions will not be the same from dataset to dataset. However, we will look at being able to provide DRL scores per configuration, in the latter parts of this Task.

There are **no specific datasets** required to assess the quality, nor specific modalities. Should there be more than 1 device type we will initially average over the values, and if a modality is missing, their contributions or weights are simply reassigned to the existing modalities.

Data sources in autonomous driving

From a dataset point of view the simplest view of the project is:

Data in ROADVIEW project = VISIBLE Camera data +
Thermal Cameras +
LiDAR data +
RADAR data +
IMU sensor positioning data +
GNSS data

To be clear, the project is large, complex and has many facets, weather modelling, testing, system integration, testing and finally working demonstrations using real trucks. That said, a vehicle that uses sensors to navigate, avoid and find its destination with respect to the road and weather conditions will be make heavy use of data, especially in the training phase. That is where the test vehicle performs journeys (see below) to learn of all possible situations it will encounter when licensed and driving passengers around. The goal from autonomous driving is to produce a "picture" of the environment, often called an **HD map**. In ROADVIEW this is discussed in additional detail in Tasks T5.4 and T8.3.

Part II: Sensors and data

Brief overview of the sensor types in the ROADVIEW project

To design and operate a vehicle in harsh weather environments, a number of sensors, platforms and processing is needed. Below we give a high-level overview, mostly for the laymen. A complete system requirement has been produced in WP2, the preparation of the demonstrators, a complete list can be found in WP2 system requirements document. Table 1 below gives a general overview of the sensor types and some characteristics.

Feature	LiDAR	RADAR	Camera	Ultrasonic
Primary Technology	Laser beam	Radio wave	Light	Sound wave
Range	∼200 m	∼250 m	∼200 m	∼5 m
Resolution	Good	Average	Very good	Poor
Affected by weather conditions	Yes	Yes	Yes	Yes
Affected by lighting conditions	No	No	Yes	No
Detects speed	Good	Very good	Poor	Poor
Detects distance	Good	Very good	Poor	Good
Interference susceptibility	Good	Poor	Very Good	Good
Size	Bulky	Small	Small	Small

Table 1 - Simple summary of sensors in autonomous driving
(source: An Overview of Autonomous Vehicles Sensors and Their Vulnerability to Weather Conditions. Sensor [28])

VISIBLE CAMERAS

Obviously, a key component in vehicle autonomy. Most cameras can be classified as visible or infrared (IR). VIS cameras such as monocular vision and stereo vision capture wavelengths that ranges from 400 to 780 nm, similar to human vision. They are mostly used due to their low cost, high resolution, and their capability to differentiate between colours. Combining two visible cameras with a predetermined focal distance allows stereo vision to be performed; hence, a 3D representation of the scene around the vehicle is possible. However, even in a stereoscopic vision camera system, the estimated depth accuracies are lower than the ones obtained from active range finders such as RADARs and LiDARs. Figure 3 shows two visible light cameras used in ROADVIEW Sekonix SF3324 and the Entron F008A030RM0A, images of the two, with specifications, are in the screenshots below.

Sekonix SF3324 – GMSL – RCCB

Sensor model	Onsemi AR0231
Resolution	2.3 MP - 1928 x 1208 px
Framerate	30 FPS
FOV	120°

Entron F001A028BM0A-

Sensor model	Onsemi AR0231
Resolution	2.3 MP - 1928 x 1208 px
Framerate	30 FPS
FOV	30°

Entron F008A030RM0A

Sensor model	Onsemi AR0820
Resolution	8.3 MP
Framerate	30 FPS
FOV	120°

Sekonix SF3325– GMSL – RCCB

Sensor model	Onsemi AR0231
Resolution	2.3 MP - 1928 x 1208 px
Framerate	30 FPS
FOV	60°

Figure 3 - Two visible light cameras used in the ROADVIEW project

THERMAL CAMARAS

Thermal cameras, also known as infrared (IR) cameras, detect infrared radiation, which is essentially heat that objects emit. This makes them different from typical cameras used in autonomous driving systems, which capture visible light. Thermal cameras can offer advantages over normal cameras in night vision, adverse weather, enhanced object detection, reduced false positives and improved braking, detecting living objects rather than non-living. ROADVIEW has access to an Adasky LW Infra-Red Vıper camera. A whitepaper about thermal cameras and the product is at link.

RADARS

Specifications: RADARs use radio waves in the 5-130GhZ range. RADAR is often classified into short range (~30m) and long range (>150m). The short radars use the 24 GHz ISM band from 24.0 to 24.25 GHz with a bandwidth of 250 MHz, also called as the narrowband (NB). Long-range radars (LRR) use the 77 GHz band, 76-81GHz, to provide better accuracy and resolution in a smaller package. Long range applications need directive antennas that provide a higher resolution within a more limited scanning range.

They are used for measuring the distance to, speed of other vehicles and detecting objects within a wider field of view e.g., for cross traffic alert systems. ROADVIEW has access to a Continental ARS 408-21, data sheet. According to the data sheet the ARS 408-21 and offers anti-collision protection, headway for far-field objects, has non-radar reflecting detection, can classify 120 objects per cluster and has distance and speed monitoring. ROADVIEW also has access to ZF's Pro Wave, with up to 350 m detection range, and 192 channels at 77 GHz. Figure 4, again, shows screenshots with illustrations and specifications.

Figure 4 - Two RADAR modules used in the ROADVIEW project

LIDARs

For the uninitiated, LiDAR devices use scanning lasers (invisible light @ 330ThZ, 1550 nm) to detect objects from a close distance to several hundred meters away, some in a 360-degree field of view. The units need to reliable for the harsh conditions of the road, rain, dust, and variable lighting, vibration, pollution and so on. They can produce a point-cloud ‘image’ of the scene around the vehicle. Often several are positioned to cover all areas around the vehicle. They generate a lot of information 5630 Kbytes per min (see Table 1 below), this can be lessened by using lower frame rates. In harsh weather the laser beams can become scattered or attenuated. FGI has access to the Velodyne VLS-128, released in 2017, link. It has 360 horizontal view (in the horizontal plane) and 40 degrees vertical. The range is 245 meters and has 0.11 minimum degrees angular. It generates 4.8 million points per second. Carissma has access to i) InnovizOne, released in 2023, has 0.1x0.1 angular resolution, 10 or 15 frame rate, 1 to 250 detection range, 115x24 degree field of view and is ISO 26262 compliant, ii) an Ouster OS1, with a 90–200-meter range, 45-degree vertical field of view, 128 channels and 5.2 million points per second. Figure 5 shows the 2 Velodyne sensors.

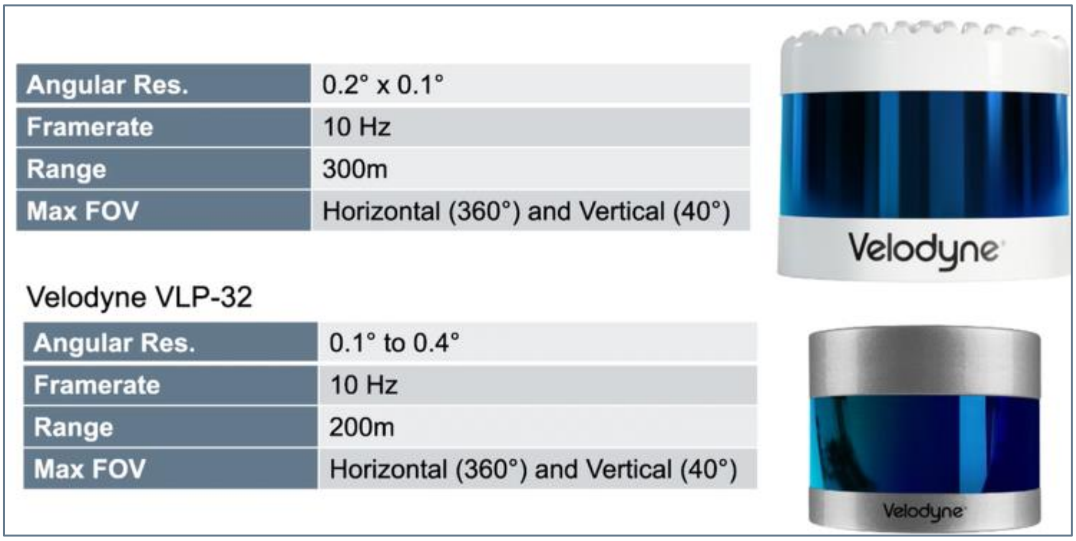

Figure 5 - Two LiDAR modules used in the ROADVIEW project

Ultrasound

Ultrasonic sensors are used for close distance obstacle recognition such as parking assistance. The reach of the sensors is typically limited to 50cm, and they are not giving you any obstacle information than the distance. And even given that, they are not as accurate as the other sensors. Pricewise, ultrasonic sensors are very low cost and in nearly every car today as parking assistance. Due to not getting any obstacle information from the ultrasonic sensor, it is not possible to classify an obstacle. Therefore, ultrasonic sensors are not taken into consideration in ROADVIEW. Credit for this information goes to Joachim Glass at Konrad-Technologies (KO).

Internal Measurement Units (IMU)

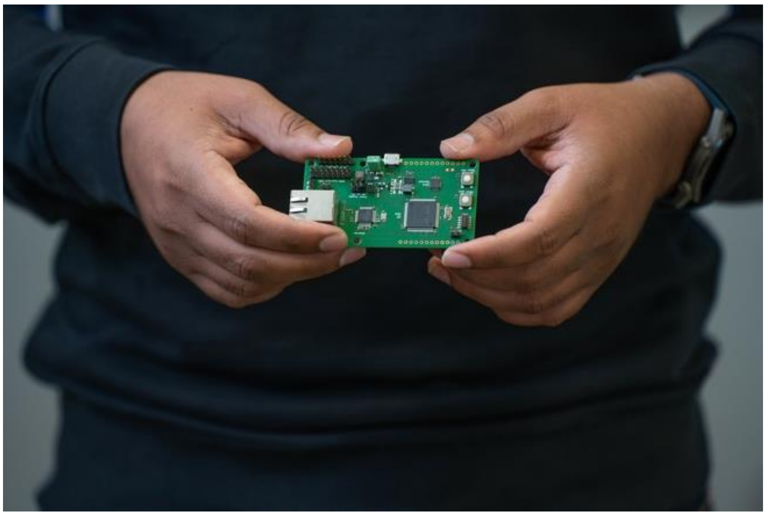

Figure 6 - An IMU unit

The Inertial Measurement Unit or IMU consists of two sensors: An accelerometer and gyroscope. The accelerometer measures the linear acceleration whereas the gyroscope measures the angular velocity. With a known starting location and precise acceleration measurements, the IMU provides information on current vehicle location and orientation. The IMU is the only sensor technology that is independent of any information from the visual or radio spectrum. The performance of IMU is limited only by the accuracy of acceleration measurements of the sensor itself. Unlike sensors such as cameras and LiDAR's, an IMU can be installed to a shielded container, deep into the vehicle's chassis. ROADVIEW has access to several IMUs. Figure 6 shows an IMU used by Sensible4, Finland.

GNSS

GNSS is the most widely used technology for providing accurate position information on the surface of the earth. The best-known GNSS system is the Global Positioning System (GPS), which is a U.S. owned utility that provides users with positioning, navigation, and timing (PNT) services. The operating principal the ability of the receiver to locate at least four satellites, calculate the distance to each & identify the receiver location using trilateration. GNSS signals suffer from several errors that degrade the accuracy, such as: (1) timing errors due to differences between the satellite atomic clock and the receiver quartz clock, (2) signal delays due to propagating through the ionosphere and troposphere, (3) multipath effect, and (4) satellite orbit uncertainties. To improve the accuracy of current positioning systems on vehicles, data from satellites are merged with data from other vehicle sensors to achieve reliable position information. In the ROADVIEW project we use Novatel Span-IGM-A1 and the XSenseMTI-710-2A8G4 GNSS sensors.

Sensor	Used in / Required for ROADVIEW Innovation	Parameter	ROADVIEW minimum sensor requirement
RGB Camera	Data Filtering Low-Level Sensor Fusion Object Detection Free Space Detection Weather-Type Estimation Slipperiness Estimation Visibility Estimation	Number of Sensors	>= 3
		Field of View	TBD
		Resolution	>= 2MP
		Framerate (Hz)	>= 10
		Sensor Model	RGB Camera
		Placement	Front, (Rear optional), both Sides
LiDAR	Data Filtering Low-Level Sensor Fusion Object Detection Free Space Detection Weather-Type Estimation Slipperiness Estimation Visibility Estimation HD Mapping Localization with HD Mapping	Number of Sensors	>= 1
		Field of View	360° horizontal >= 40° vertical
		Resolution	128 bin
		Framerate (Hz)	>= 10
		Measurement Range (m)	>= 80
		Sensor Model	No Requirement
		Placement	Top (to minimize obstacles)
RADAR	Data Filtering Low-Level Sensor Fusion Object Detection Free Space Detection Weather-Type Estimation Visibility Estimation	Number of Sensors	>= 1
		Field of View	TBD
		Resolution	TBD
		Framerate (Hz)	>= 10
		Measurement Range (m)	>= 100
		Sensor Model	No Requirement
		Placement	Front
Thermal Camera	Low-Level Sensor Fusion Object Detection Free Space Detection Weather-Type Estimation Slipperiness Estimation	Number of Sensors	>= 1
		Field of View	TBD
		Resolution	TBD
		Framerate (Hz)	>= 10
		Sensor Model	No Requirement
		Placement	Front
GNSS	HD Mapping Localization with HD Mapping	Number of Sensors	1
		Framerate (Hz)	>= 10
		Sensor Model	No Requirement
		Placement	Antenna on Top
IMU	HD Mapping Localization with HD Mapping	Number of Sensors	1
		Framerate (Hz)	>= 100
		Sensor Model	No Requirement

		Placement	No Requirement

Table 2- Minimum requirements on ROADVIEW sensor set, being defined in Task 4.1. The Data Readiness Level will be dependent on the setup, this is one example of what a setup will be.

HD Mapping

An HD (High Definition) map is a detailed and highly accurate digital representation of the real-world environment, primarily developed for autonomous driving systems. These maps go beyond the traditional navigation maps that we use in our smartphones or cars. Here's what distinguishes HD maps from standard maps:

- Detail and Precision: HD maps provide centimetre-level precision, enabling self-driving cars to understand their surroundings better and make informed decisions.
- Layers: Unlike standard maps that might only provide roads and points of interest, HD maps come with multiple layers of information, including:
- Road Profile: This includes details like road curvature, gradient, and width.
- Lane Information: It offers specifics about each lane, its boundaries, type (e.g., turning lane, straight lane), and associated rules.
- Traffic Signs & Signals: HD maps will have exact positions of all traffic signs, signals, and other regulatory markers.
- Infrastructure Details: This may include crosswalks, barriers, guardrails, pedestrian areas, and more.

Dynamic Updates: Given the rapidly changing nature of roads due to construction, accidents, or other events, it's crucial for HD maps to be frequently and dynamically updated. Some systems aim to update in near real-time. 3D Representation: While many standard maps offer 3D views for a better user experience, HD maps can provide a true 3D representation, accounting for elevation changes and including structures like bridges, tunnels, and buildings. Sensors & Integration: HD maps are typically developed considering the suite of sensors (like LIDAR, radar, cameras) on autonomous vehicles. The data from these sensors can be cross-referenced with HD maps for tasks such as precise localisation. HD maps play a crucial role in making autonomous driving safer and more reliable. By giving vehicles, a comprehensive understanding of their environment, they help ensure that the vehicle can handle complex driving scenarios even if sensors face temporary obstructions or difficulties.

A ROADVIEW data pipeline

Requirements (WP2) ⇒ Sensor capture (WP4) ⇒ Format wrangling (WP4) ⇒ Validation checks (WP4) ⇒ Fused data (WP4) ⇒ Model training (WP5) ⇒ Object detection (WP5)

Data readiness, or quality can be defined for the online or offline case. In the online case could be a situation that was not anticipated or not seen before causing the vehicle to make a poor decision (the classic Über vehicle crash in the US, unfortunately killing a pedestrian). The offline case is pertinent to highlighting of internal practices, where the software system can be "debugged" from a data point of view. Should object detection fail we can look at the causes due to data issues. Note, the algorithms themselves could be issue. These should be tested on a 'perfect dataset' possibly machine generated (see below).

Data quality & autonomous vehicles: a short (& selected) state of the art

Since autonomous vehicles are essentially 'computers on wheels' data gathering, processing & navigating there is a substantial amount of material. Often citing each sensor type, and increasingly more often, fused data. The photometric society of America produced a paper at the link about Geometric Inter-Swath Accuracy and Quality of LiDAR Data, which is not related to driving but an interesting take on LiDAR data quality. Data quality issues can be found in [1-3]. Within the consortium [4-6] cover vehicle segmentation in LiDAR point clouds. Adverse weather autonomous driving is in [7-13]. Each public datasets have its own set of publications, where nuScenes and mmDetection3D attempt to summarise the others, NuScences 'corruptions' a method is close to our work [20]. Around ML pipelines references [15-18] are relevant, where we point to [15] as a lightweight approach. [14] is video quality estimation done by video quality experts, rather than vehicle developers or ML specialists.

Data quality for autonomous vehicles: image quality

A key aspect in this project is the use of multiple sensors in adverse or poor weather conditions. It is anticipated that driving in conditions with non-differentiating backgrounds, water particles in the air (attenuating sensors) make it more difficult for the vehicle to assess surrounding obstacles and obstructions. Visible image quality issues come from lack of sharpness due to focus or dirty sensors, blurriness from lack of focus, confusing ranges and even the speed of the measurement vehicle itself. Obviously over exposed images from bright or low sun / lights as well as reflectance from piles of snow (see citations). If video processing is done on a sequence of images, then single frame 'images' can be factor especially over several ones, also due to compression of both images and video. Some form of downscaling is needed, either in frame rates or redundant coding. Below in Table 3 are a couple of image quality assessment tools. Future work will experiment with additional image assessment.

Image Attribute	Source Lang.	Comment	GitHub Stars	Forks	Links / Notes.
Blurriness	Python	Provides a quick & accurate method for scoring blurriness. See Figure 8.	277	1	https://github.com/WillBrennan/BlurDetection2
Multiple	Python	PyTorch Image Quality (PIQ) is a collection of measures and metrics for image quality assessment. The library contains a set of measures and metrics that is continually getting extended	1100	107	https://github.com/photosynthesis-team/piq
Blurriness	Python	Attempts to judge if an image is blurred or not using a score.	-	-	Own development.
Exposure	Python	No reference image sharpness assessment based on local phase coherence.	40	7	https://github.com/elejke/awesome-defocus-detection

Sharpness1	Python	Sharpness methods.	31	7	https://github.com/topics/image-sharpness
Sharpness2	Python	Sharpness detection	-	-	Library calculates the variance of the Laplacian for each greyscale image.
Brightness	Python	Mean of the pixel values for the greyscale images Over and underexposure	-	-	Own development.
Noise	Python	Library to calculate Shannon entropy of each image skimage.	Many and varied.		https://github.com/topics/skimage
LiDAR Attribute	**Source Lang.**	**Comment**	**GitHub Stars**	**Forks**	**Links / Notes.**
Point-Cloud Noise	Python	A complete toolbox from Open mmDetection3D.	27.9K	9.2K	https://github.com/open-mmlab/mmdetection Corruption / noise in point clouds. See Figure 18.
Point-Cloud Noise	Python	OpenPCDet is a clear, simple, self-contained open-source project for LiDAR-based 3D object detection.	4.3K	1.2K	https://github.com/open-mmlab/OpenPCDet
Point-Cloud Noise	Python	MultiCorrupt: Weather effects in the point clouds [29].			https://github.com/ika-rwth-aachen/MultiCorrupt

Table 3 - Image & LiDAR quality assessment with open-source repos

Still versus motion quality assessment, the quality assessment for video has some options, `ffmpeg`, the open-source audio and video codec, probably is the best-known for video, notably in cross platform players such as VLC for decoding (usually) and HandBrake (coding). It also includes three quality metrics:

- Peak Signal-to-Noise Ratio (PSNR), which measures the difference between the original and compressed videos. A higher PSNR generally indicates that the reconstruction is of higher quality.
- Video Multi-Method Assessment Fusion (VMAF) developed by Netflix, VMAF is a perceptual quality metric that considers both human vision system models and machine learning models to provide a more accurate quality score.
- Video Structural Similarity Index (SSIM) is another metric that evaluates the perceptual difference between two videos. A value closer to 1 means the videos are more similar.

We can use `ffmpeg` as the external tool videos, as it has `ffprobe` and `ffplay` as sister tools to play, examine and look at the coding, these tools help when dealing with many files. In the VMAF case, a separate codebase and tool is available. Note, however a reference or “best quality” is needed, to make the comparison. Note others exist [14], as well as separate ones for LiDAR [20, 21]. A no reference video quality paper is in [26]. If videos are recoded, or resized or even format changed (mkv, mp4, mov, wmv) one should use the quality comparison tools above. An interesting, and yet unexplored option is to rate images from good weather conditions. In other words, ideal conditions against poor/harsh/adverse conditions, again in other words again, the reference image is the good weather, and the degraded image is in poor weather conditions.

Data Readiness Levels: The concept

Data Readiness Level (DRL): In analogy with the concept of technology readiness levels [1], data readiness level assessment is a method for characterising data readiness for deployment. The overall data quality builds on Lawrence's data readiness levels (DRL) concept [2] and is a significant part of project planning and development. It is not by chance that up to 80% of the total project time is spent on pre-processing data, basically following the Pareto principle [1]. The principle states "it takes 20% of the time to do 80% of the work, and 80% of the time to do the remaining 20%", even simpler stated, "one can do the large parts, but the fiddling consumes a lot of time (and effort)".

The main challenges of constructing meaningful data readiness levels are: i) assigning a single DRL (1–9) to large, complex datasets, often with imperfections [2], such as missing values, inaccuracies, and incomplete readings; ii) different readiness suggests different implications to different users, since data is often context sensitive; iii) some data consumers may have methods to handle imperfections, for example, in the missing data case, methods may, or may not, have been coded to handle missing values. Depending on the upstream capabilities, the DRL may be inaccurate [3,4]; and iv) production of quality sensor datasets (real car readings) are currently available for use.

One way to see data readiness is the values 1-9 indicate a quantitative measure of the time / effort / cost to repair / replace produce new values. Lower values are in principle easier to fix than those higher up the scale. Note, if errors in the lower end will propagate through the data pipeline causing issues, and probably harder to detect. This is because transformations (scaling, fusing, ML algorithms etc.) are applied and debugging the ML pipeline is more time consuming. Indeed, WP5 deals with making the ML pipeline as transparent as possible.

Lawrence's initial concept defines data readiness in three different bands—A (utility), B (validity), and C (accessibility)—depending on the knowledge and understanding of the available data and their usefulness for a given objective. We go back to the 9 levels present in TRLs and use them as below in Table 4 .

Figure 7- Lawrence's data quality bands

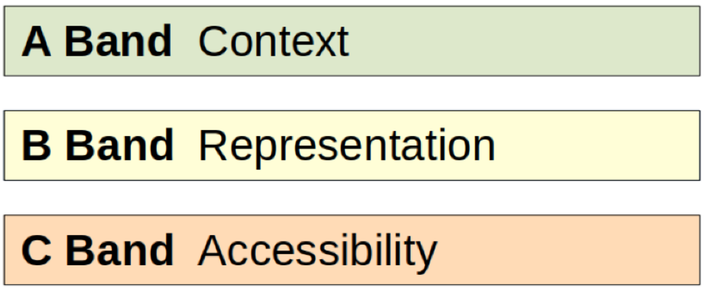

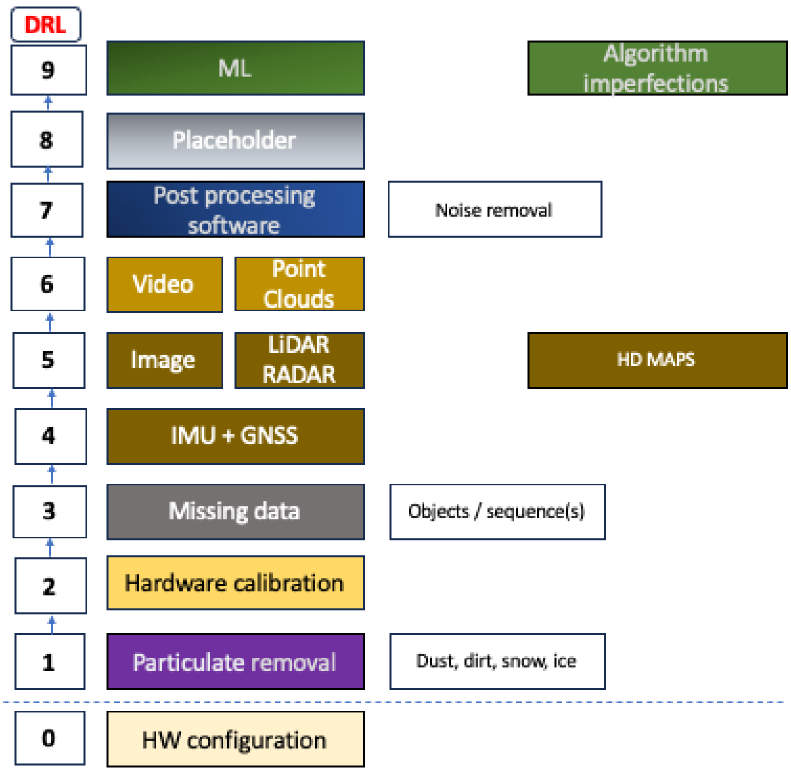

Figure 8 - Illustration of the DRLs for ROADVIEW

Level	Type		Description	Example / Notes
9 (highest)	ML		Successful object recognition in harsh weather in a live driving situation.	A vehicle driving around rural and urban environments.
8	ML		Successful computer object recognition in harsh weather conditions.	An actual vehicle is not involved, rather in a lab, (offline) test. Probably
7	ML		Insufficient / incorrect data for training.	
6	Video	Video frames	Quality issues	Video quality experts' group. VISIBLE and thermal different.
5	Video	Video frame	Missing values	Gaps in frames causes problems.
4	Images	Single image	Brightness, contrast, blurriness	

3	Images	Single image	Missing streamed values	Gaps in the data sequence (Missing frames in VisualizePixelwiseFusionImages., FGI)
2	Data	Single image	Incorrect data	Confusing / wrong information. Fields used in inconsistent manners Data repeated in non-normalised DB schemas.
1 (lowest)	Data		Formatting	Mismatches (decimal separators) points, km / mph

Table 4 - Specific Data Readiness Levels for ROADVIEW. The light-yellow shaded cells indicate Domain specific data issues whilst the light blue cells data agnostic processing.

Data source I: The Finnish Geological Institute (FGI) training data

We take a little deeper dive into a dataset available in the ROADVIEW project. It was kindly made available by the Finish Geological Institute; we list some of the attributes of the data made available. The scope of the description is quite high-level formats, visualisation to the lower-level formatting and how to unpack and use the four sensor types. Table 5 shows the sensor, the recorded attribute and format and notes on the data.

FGI sensors, formats, and files

Sensor	Format	Notes
Camera	PNG	Anonymised (Blurred faces and number plates). Unrestricted license compressed format.
Thermal Camera	TIFF	Raw, unrestricted license format (16 bit)
LiDAR	RG	Binary packed range and reflection data
Road weather	RW	Binary packed road conditions data (see below)
RADAR	-	Not available from FGI (yet)

Table 5 - Data formats from FGI dataset

One aspect of FGI's data is the use of binary weather formats, binary structured files. FGI use Python's:

```
Struct.unpack_from(format, data, offset)
```

Function to pull out the file's data. This we needed to generalise for reading other datasets, in other formats such as ROSBAGs. Waymo uses HDF5 (see below).

Road conditions

In the FGI dataset, the road conditions were categorised as in Table 6 . Bool represents Boolean, if available and measured or not.

Attribute	Format
Surface Temperature	Bool
State	Bool
Water	Bool
Grip	Bool
Ice	Bool
Snow	Bool
EN15518 State	Bool
Air Temperature	Bool
RH	Bool
Dew Point	Bool
Frost Point	Bool
Data Warning	Bool
Data Error	Bool
Unit Status	Bool
Error Bits	Bool

Table 6 - Road Weather, conditions & data fields in the FGI Dataset

A master's thesis on the topic has recently been produced, which includes slipperiness and grip [27].

Urban and rural driving data

The Finish Geospatial Institute (FGI) has produced a data set, shown in Table 5 . It was used to generate the video shown in link. The whole data set size was about 34 Gigabytes for a 10-minute drive. 8 below shows the data rates and message sizes computed in kB / s and in MB / minute. Content-wise the data is divided into a Rural and Urban drives, 1810 and 3027 in PNG, TIFF images and RW 'format' (see above) respectively.

Rural	Sizes	# Files	Urban	Sizes	# Files
Camera anonymised	2.4G	1810	Camera anonymised	4.7G	3027
LiDAR	913M	1810	LiDAR	1.5G	3027
Thermal camera	6.7G	1810	Thermal camera	39M	3027
Road Weather	21M	1810	Road Weather	39M	3027

Table 7 - Two autonomous driving datasets from FGI (https://www.maanmittauslaitos.fi/en/research)

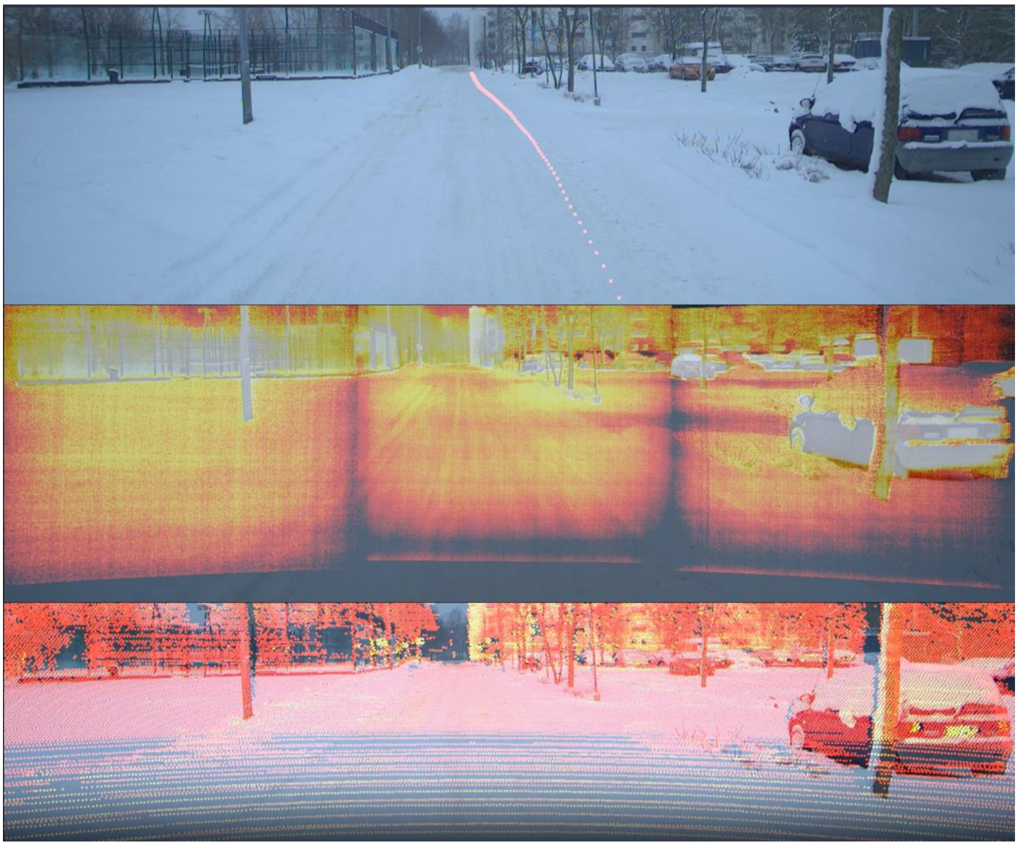

Figure 9 - Example of FGIs Urban 3 data sources fused, Video link (5fps).

The output of a processed urban dataset is a 250 MB video around 5 minutes long at 10 frames per second. The rural one is 60 MB around 3 minutes also at 10 frames per second. More on rates and sizes in Table 7 above. Note this is without RADAR data and does not say anything about training, validating an AI model. The start frames can also be specified and in the rural case this is offset by 300 'frames' as the capturing started early. Table 8 shows the sizes and times to fuse data (offline) the two datasets. Processing was done on a 2021 MacBook Pro (M1 Max CPU).

Environment	Output video (@ 10 frames / sec)	Fusion Time
Rural	60 MB	30 mins
Urban	250 MB	50 mins

Table 8 - Output characteristics using the FGI dataset above.

In terms of performance issues, Python will produce binary files *.pyc rather than text to make the interpreted code into a binary. For large programs (rather than data) this can speed up execution. For optimising the code, one can use `timeit` to find functions or sequences that consume time. Alternatives are `cProfile` and very recently, `Scalene` from Umass, Amherst, USA (Aug 28th, 2023). Table 9 below gives the breakdown per sensor and how the data streams are reduced by down sampling.

Sensor	Data per message (kB)	Freq. (Hz)	Data rate (kB/s)	MB/min	Data rate in 5Hz low sampled data (kB/s)	MB/min
Novatel GNSS + IMU positioning	0,09	205,0	19	1,10	19	1,10
Front colour camera	1399	10,0	13984	819	6993	410
Thermal camera centre	241	60,0	14459	847	1204	71
Thermal camera left	240	60,0	14407	844	1200	70
Thermal camera right	235	60,1	14126	828	1176	69
Vaisala MD30	0,04	40,0	2	0,09	2	0,09
Velodyne VLS128 LiDAR	617	9,1	5630	330	3087	181
SUM			**62626**	**3669**	**13680**	**802**

Table 9 - Sensor sizes and rates from the FGI dataset.

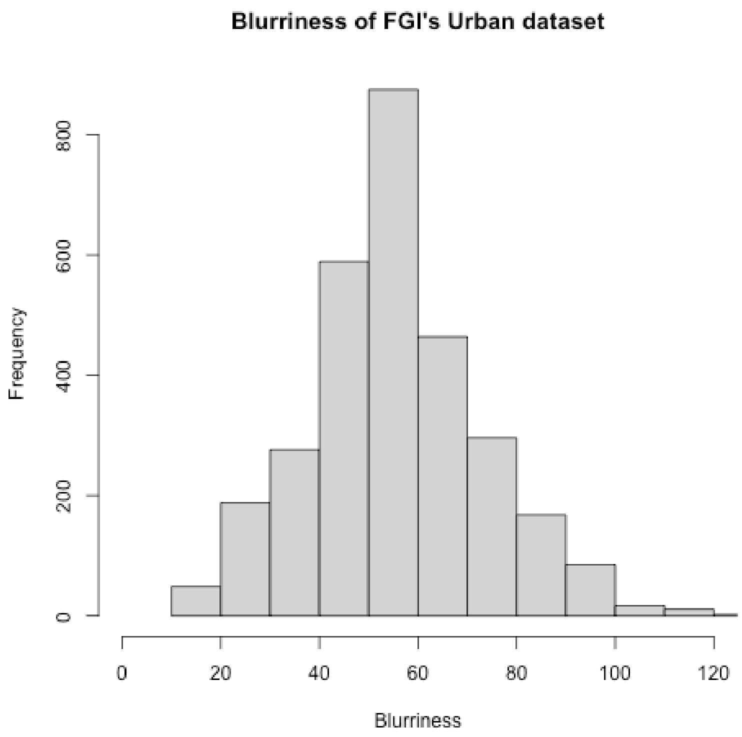

Figure 10 - Blurriness obtained by taking the variation of the Laplacian of the 3000 images in the FGi's dataset. Image sizes are scaled to HD. Higher values are less blurry.

As can be noted, thermal cameras have a lot higher data rate than other sensors (60Hz vs 10Hz). Now there is 3.6 GB of raw data per minute. If only the frames which are used to generate the sensor-fused frames at 5Hz are listed, the data rates decrease significantly. Computed is a low sampled example in the two last columns, where only 5 samples per second have been taken from all cameras and the LiDAR. In this low-sampled version one minute of raw data is 802 MB. The project will take different data sets under various weather conditions. Using a low sampled would take a few 100's GBs, but if full resolution is needed, an estimate is > 1TB.

In the initial readiness assessment, RISE introduced the concept of DRLs. Data plays a fundamental role in the ROADVIEW project, making the completeness, correctness and interoperability of data are utmost important aspects of the data processing pipeline. RISE investigated the multimodal sensors that are used in ROADVIEW, summarizing the physical nature of each sensor, but focusing on camera image quality for the initial DRL definition in D4.5. As shown in Figure 8, from the lower DRL levels, **DRL 1** is about the sensors and perturbations therefrom. We know snow, ice, dust, and dirt inhibit signals from the LIDAR and RADAR. This is very much the focus of other tasks, with a slight emphasis on RADAR. Calibration at **DRL 2** is an important aspect, calibration, which we separate into 2 types:

Intrinsic calibration
Intrinsic calibration is basically sensor calibration, lens alignment, which can result in distorted / rotated images and point clouds, the resulting matrices contain wrong values with respect to reality, often abbreviated to NAs. The abbreviation can stand for Not Applicable, Not Available or Not Assessed. In the case of LiDAR, the laser sweeps should be within a range and intrinsic calibration is done to a known point cloud, physically adjusting the LiDARs. Often RGB cameras are seen as the trusted modality. Intrinsic calibration is only done once, rechecks and calibration are redone if something breaks during trials.

Extrinsic calibration
It is important to position the sensors with respect to each other. In vehicles with sensors, often the RGB camera is seen as the 'base' sensor or modality. One reason is that one can infer 'where' the camera is on the vehicle using internal and external imaging solutions. Speciality software takes pictures outside of the vehicle and inside the vehicle to calculate the position of the RGB cameras. The Radio Cross Section (RCS) of a target is the equivalent area as seen by a RADAR. It is the fictitious area intercepting that amount of power which, when scattered equally in all directions, produces an echo at the RADAR equal to that from the target.

At **DRL 3** we deal with missing data, as when streamed from sensors in real time (even when buffered) gaps in the sequence can occur. Data rates exceed capacity, sensors do not always pick up information, noise in the data overrides the signal and so on. The ratio of signal to noise is a key factor in any sensing situation. Theoretical results and extensive practice and calibration can to some degree predict the performance of a sensor, or modality in autonomous driving. However, in practice, and real driving situations it is more difficult to ascertain the signal and noise ratios. In the dataset from FGI a few values are missing, see Figure 11. Two representations of 'sequence gaps' in the FGI dataset. The leftmost by differencing, the rightmost by absolute timestamps below.

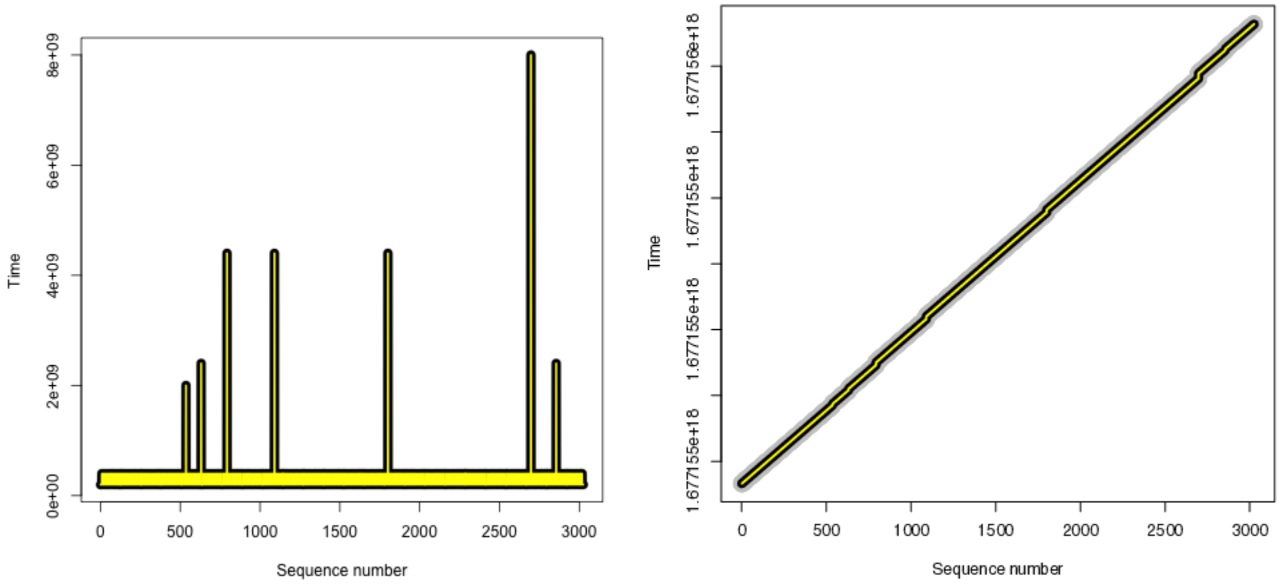

Figure 11 - Two representations of 'sequence gaps' in the FGI dataset. The leftmost by differencing, the rightmost by absolute timestamps.

At **DRL 4** we have the positioning and inertia measurement units. For now, we have them at the same level. This could be for discussion later, but knowing where we are (GNSS) and how the vehicle accelerates, brakes, or vibrates is important in post-processing. Repeated runs along the same route in different conditions produces repeatability but also changing, a limited, but known, number of parameters. Indeed, in ROADVIEW runs are done when the weather has changed. An extension is to create the weather artificially, as done in the THI and Cerema test site in Germany and France.

At **DRL 5** we have images and the LIDAR / RADAR. They are different modalities, but in essence are similar in terms of sensing. We look at several different open-source tools to assess them.

At **DRL 6** we have the collated 'data'. That is video and LiDAR scans, processed (frame rates, resolutions) single images combined into videos and LiDAR point clouds, spatially and temporally collated. We look at the video quality as well in this Task, using `vmaf` (Table 61) which the rest of the project does not do. LiDAR point clouds are heavily discussed in this report. With respect to data quality, we have access to the FGI & THI datasets and are evaluating those, and we can introduce weather effects into public datasets and evaluate them, this is done with tools such as MultiCorrupt [29] and Robo3D.

At **DRL 7** we have the processing of the point clouds, removing noise from the point clouds to make object recognition effective, this is being done in other parts of the project.

We have left one **DRL 8** free as some processing in the ML pipeline, that can affect the final quality. Research is being done to rasterise point clouds, make the point clouds more sense or using multiple sources to enhance quality. Therefore, we leave **DRL 8** for each advancement.

DRL 9 is strictly about the object detection, in free space or in scenes that is coupled to the algorithms themselves. We have implicitly assumed that the algorithms are perfect and only data is a detractor, which is strictly not true. The algorithm and often the algorithm and the data it is trained, validated, and tested upon can have significant impacts on the results in AD. This is why different techniques are evaluated, often on several datasets, and with quite different outcomes. The `mmDetection` suite tests their algorithms on different datasets showing a fair difference. Therefore, the algorithms themselves need to be considered, and this is being done in WP5.

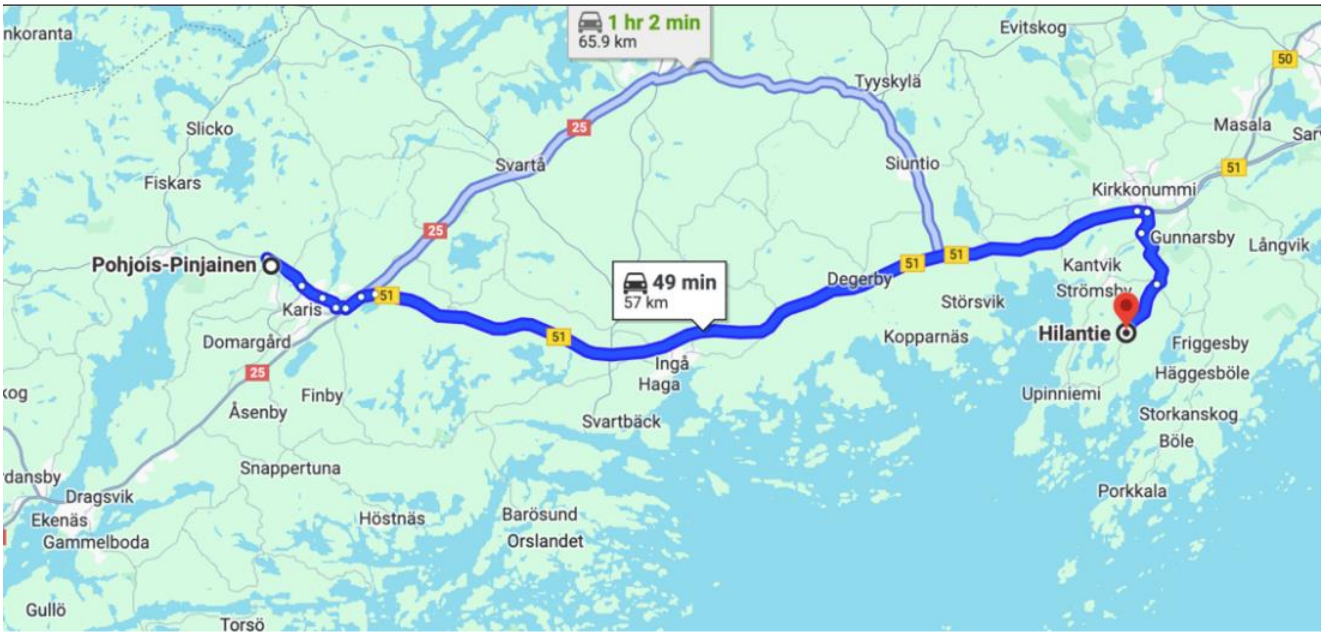

Figure 12 - Otaniemi. A rural route between Hilantie-Pohjoiseen in southern Finland.

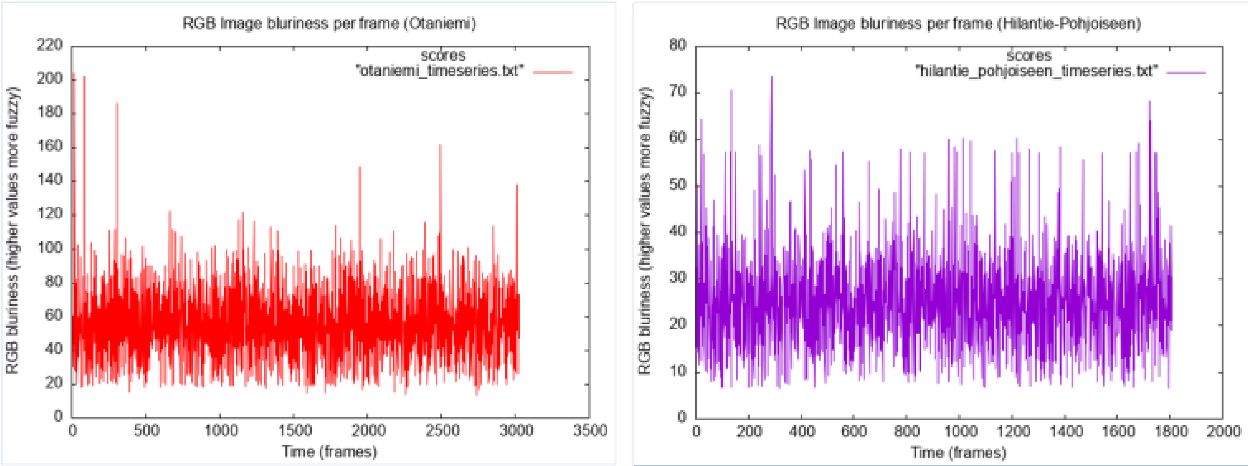

Figure 13 - **Time series of blurriness. Left: recording in an urban environment. Right: recording in rural area / transit.**

Linear combinations of sensor inputs

The main fulcrum of this work is how the DRL is calculated. For the moment, it is a linear combination of the inputs from the sensors. Should one be missing, RADAR for example, its contribution is distributed amongst the others. Similarly, so with multiple numbers of the same, for example 4 cameras. For forward and backward sensors, we will (for now) exclusively use the forward ones. Side ones will be given 25% max. However, this is a heuristic value and will be explored further.

It is important to reiterate that the Data Readiness or Quality is really for off-line use, to ensure the data is of sufficient quality. It is assumed that the algorithms are 'perfect' which is unreasonable, but there are many algorithms and measures to determine how good the algorithm is. Estimating the data quality known the performance (in %) of algorithms is possible, but over complicates the issue and might place emphasis on the wrong place.

An alternative would be start from the other end and say when the object detection is correct in the training data, then the data is perfect and for each 'imperfection' in the image correlate this with the classification success. Or a misclassification is back traced to an image that is not clear enough. There is also the issue of the training and validation sets, how they are selected. In the standard image sets this is done, however in our internal datasets this is being assessed.

DRL = VISIBLE imaging + Thermal imaging + LiDAR imaging + RADAR imaging

Implementing DRLs for ROADVIEW

In the first incarnation we used the FGI implementation, which we first refactored, from `plotData_example.py` source code with four functions:

```
def ReadRWImage(filename, out_datas=None, ret_image=False):

def ReadRangeImage(filename, ret_image=False):

def ExpandPixels(image, amount):
```

```
def VisualizePixelwiseFusionImages(main_folder, start_index=0, show_fused=False):
```

Consisting of 479 lines. RISE increased the modularity from the 4 functions to 13 functions, and no. of code lines 601. As well as to understand the code better, reduce repeated code, make hooks for reading other data sources, as well as the DRL functions. The code is available from RISE Git Lab. RISE is looking at ice cloud shared notebooks, for now we will share code at the link above internally and use ROADVIEWs GitHub repository.

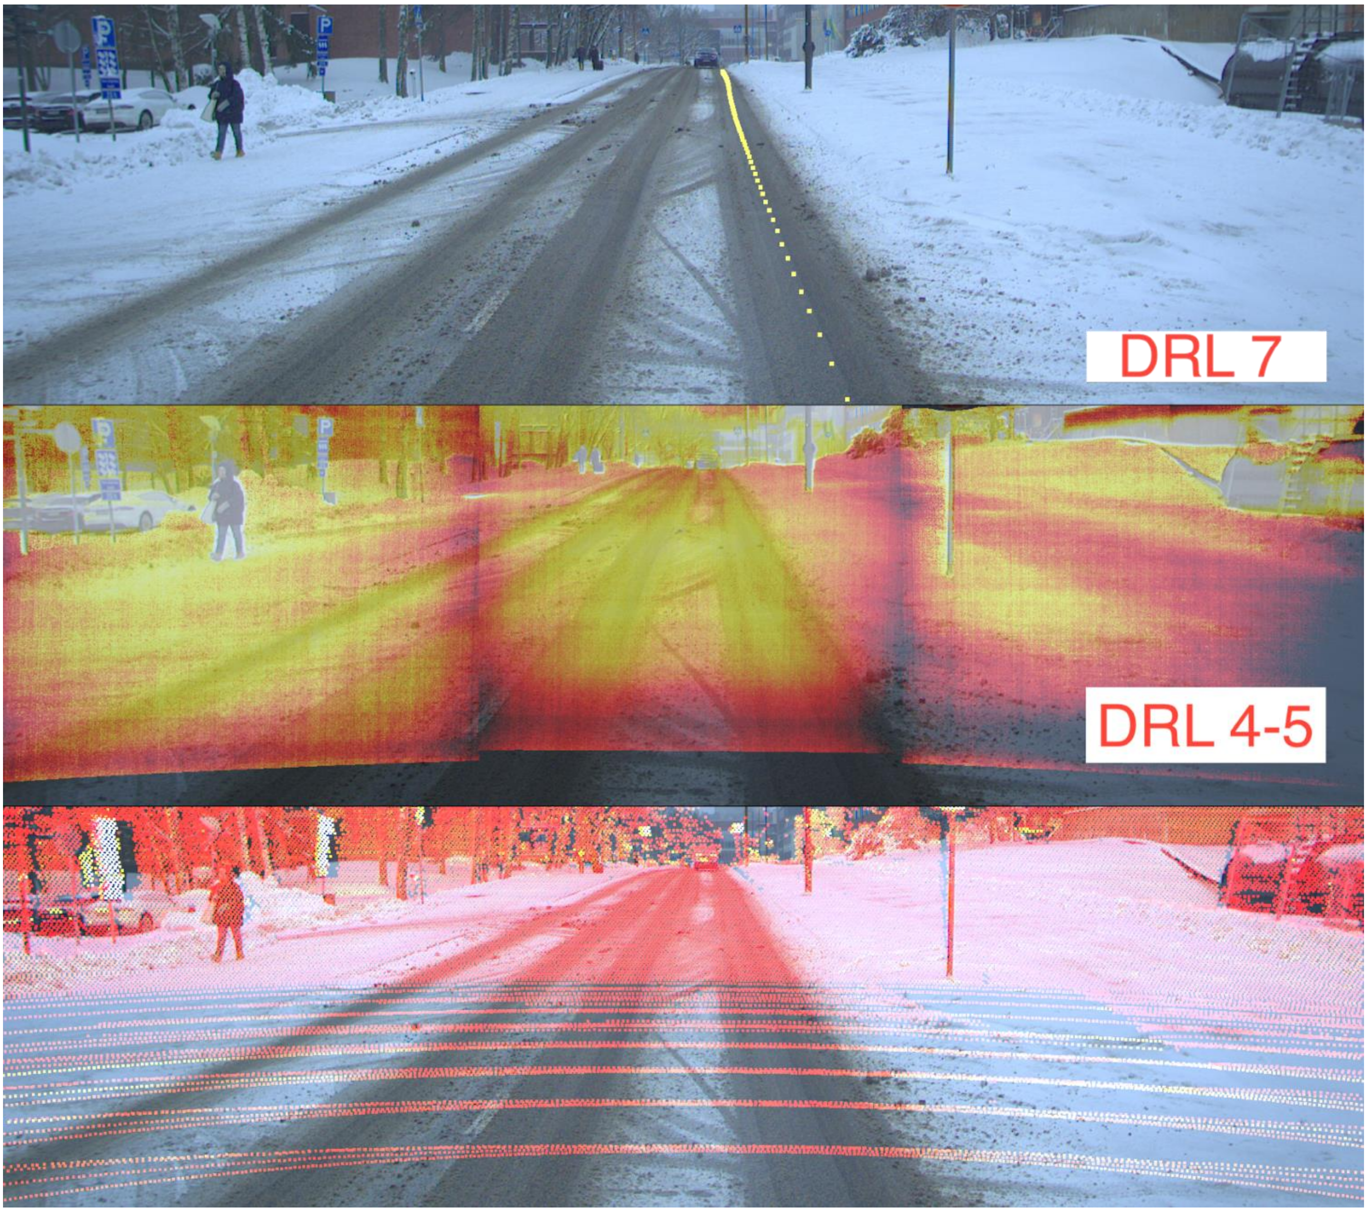

Figure 14 - DRLs applied to the images from FGI urban journey (credit: The Finish Geospatial Institute, fgi.fi). Top: Simple RGB image. Mid: thermal camera image projected into the RGB camera image. Bottom: LiDAR point cloud projected into RGB camera image.

Finally, for the FGI dataset, Figure 14 shows one representative frame of the dataset, showing the RGB camera as well as thermal counterpart and LiDAR point cloud data, all projected into the RGB camera image. The RGB camera image was rated with DRL 7 while the thermal camera frame only got up to 4 – 5. This is mainly due to the edges of the thermal camera image, simply caused by the working principle of the sensor itself. These edges do cause issues in the image quality evaluation though, hence less DRL rating. Lastly the LiDAR point cloud is yet to be evaluated since the methodology for point cloud rating is not yet finalized. The projected point cloud into the image plane in the bottom image of Figure 12 leads to the question of rating extrinsic calibration, which needs to be further discussed also. Ongoing work will look at the dataset from THI and VTT which has been received by RISE. Together with HH, the LiDAR point cloud data will be projected into spherical coordinate system to generate image representations for assessment. RISE also plans to apply the DRL concept to at least 1 public dataset.

Figure 8 showed the ROADVIEW reference architecture together with the DRL levels. The data readiness level is based on the perception module. The DRLs follow the data flow as indicated in the DRL architecture to the right. Data processing steps that occur in practice that influence the data quality occur in the DRL data processing pipeline are shown in the right figure. An example is dealing with missing data, that is a fact in real time sensing and importantly can detract from the overall quality, furthermore, affects the steps upstream. Therefore, there is no direct 1-1 mapping from the reference architecture to the Data Readiness Levels but is close (and should be).

With the modalities with the datasets we have, FGI is all minus RADAR and THI is all modalities.

The next steps for Task 4.3:

- Continue working with datasets wrt. THI, VTT
- Incorporate RADAR data from RISE
- Look at missing data impacts from VTT
- Investigate how calibration affects DRLs (upwards)
- Look at rasterising point clouds to images

Main outcomes
Initial data readiness levels with respect to the FGI dataset. Developed a software framework for DRLs and open datasets. Evaluated 5 image and 2 LiDAR point cloud datasets.

Data annotations

Whilst most of this report is about data quality, annotating data is a significant issue. In machine learning, data annotation is the process of labelling data. This labelled data is then used to train supervised learning models. Data annotation is a crucial step in many machine learning projects, especially in projects such as ROADVIEW: vision, scene detection. Obviously, the quality and accuracy of the annotated data affects the performance of the resulting machine learning models. Within ROADVIEW we use:

- Image Annotation: This involves marking various objects within images. Types of image annotations include:
- Bounding boxes: Drawing rectangles around objects of interest.
- Semantic segmentation: Labelling each pixel of an image with a class label.
- Polygonal segmentation: Drawing polygons around objects, especially those with irregular shapes.
- Key point annotation: Marking specific points of interest on objects, often used for pose estimation.
- Named Entity Recognition (NER): Labelling words or sequences of words as specific entities like names, locations, or dates.
- Video Annotation: Labelling objects or actions within video sequences. This can involve bounding boxes over time or tagging entire video clips with action labels.
- Audio Annotation: Marking sections of audio data to label different sounds, words, or other audible events.

The process of data annotation is time-consuming, requires domain expertise and ROADVIEW has a Task for this purpose. An interesting angle is active learning, where a machine learning paradigm where the model itself decides which data points should be annotated furthermore, based on where it predicts the annotation would be most valuable. This can help reduce the amount of manual annotation.

Data source II: ROADVIEW-Carissma

Below is the test track and examples of the CARISSMA outdoor test facilities. There is an acceleration zone of 210m and a dynamic area with a 60m x 70m area. The maximum speed is 100km/h. Different parts of the track are watered with different intensities. The. Images show the watering and measurement instrumentation. Initial data sets with images and point clouds have been produced. The full work will be available in Feb. 2024. Figure 15 shows the test track, 'artificial' weather conditions (rain) and a LiDAR point cloud in these conditions.

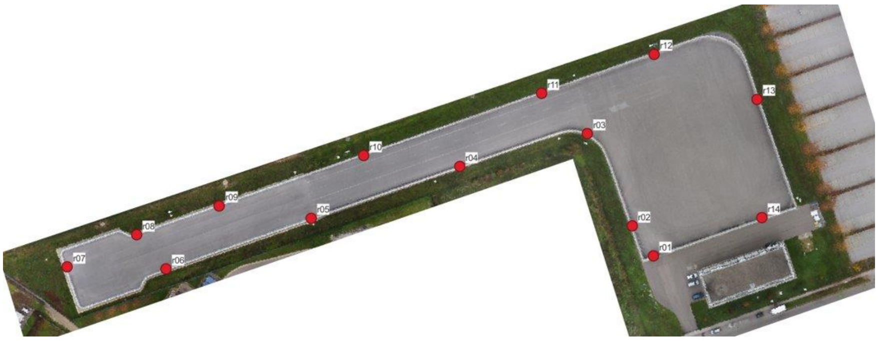

Figure 15 - Refence Dataset of measured weather characteristics, THI D3.2, WP3

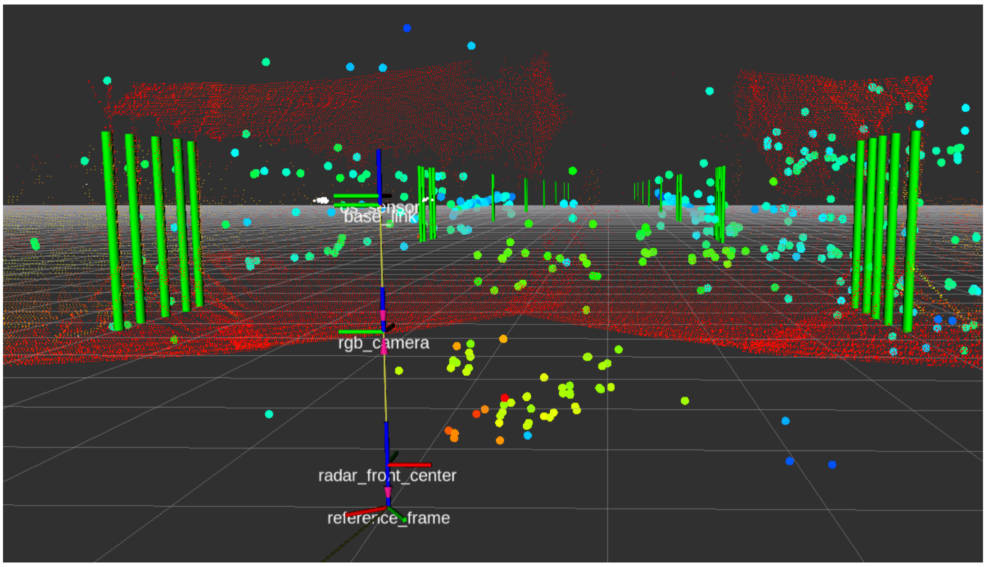

Figure 16 - Reference Dataset of measured weather characteristics, THI D3.2, WP3

The datasets were created to observe the 3 weather conditions rain, clear, and fog different conditions on different sensors. The utilised sensors in this case were VISIBLE Camera (LUCID), FLIR (Thermal) Camera, RADAR (ZF ProWave), and LiDAR (Innoviz One and Ouster OS1). This dataset has a total of 68.4 minutes of recording. The recording took place in the Carissma outdoor proving ground in Ingolstadt Germany, and in the CE proving ground in Clermont Ferrand France. Each with different intensities. The amount of rain was also measured and calibrated using three different methods, litres per square meter, drop shape and amount, and direct weather measurements.

Data Source III: Open-source data sources

Recently companies & research institutions have made their autonomous driving datasets open to the public. A Medium post found and summarised 15, as of July 2021, by Alex Nguyen. **Audi A2D2** dataset 41K labelled images with 38 features, 2.3 TB split by annotation type, semantic segmentation, 3D bounding box. **ApolloScape**, 100K street view frames, 80k LiDAR point cloud and 1000km trajectories for urban traffic. 3D tracking annotations for 113 scenes and over 324,000 unique vehicle trajectories for motion forecasting. **Berkeley DeepDrive** 100k annotated videos and 10 tasks, 1000 hours driving, 100M frames plus geographic, environmental, and weather diversity. **Cityscapes** urban street scenes in 50 German cities. Semantic, instance-wise, and dense pixel annotations for 30 classes grouped into 8 categories, 5K images with fine annotations & 20K with coarse annotations. **Comma2k19**, 33 hours of commute time recorded on highway 280 in California. 1-minute scenes captured on 20km of highway between San Jose-San Francisco. Collected using comma EONs, i.e., a road-facing camera, phone GPS, thermometers, and a 9-axis IMU. **Google Landmarks** (2018) divided into two sets of images to evaluate recognition and retrieval of human-made and natural landmarks. 2M images of 30K unique world, 2019 saw Landmarks-v2, 5M images & 200K landmarks. **KITTI Vision data**, 2012**. LeddarTech** Dataset, 2021, cameras, LiDARs, radar, IMU + full waveform a 3D solid-state flash LiDAR sensor. Contains 29k frames in 97 sequences & 1.3M 3D boxes annotated. **Level 5 Open Data** (Lyft) 55K human-labelled 3D annotated frames, surface map, and an underlying HD spatial semantic map captured by 7 cameras and 3 LiDAR sensors. **nuScenes** dataset from Boston + Singapore using a full sensor suite, 32-beam LiDAR, 6 360° cameras and radars, the dataset with 1.44M camera images capturing a diverse range of traffic situations, driving manoeuvres, and unexpected behaviours. Examples are from clear weather night-time, rain & construction zones. **Oxford Radar RobotCar** Dataset contains 100+ recordings of a route through Oxford, UK, captured over 1 year. Captures different conditions, including weather, traffic & pedestrians + construction / roadworks. **PandaSet** the 1st open-source AV dataset for academic & commercial use. Contains 48K camera images, 16K LiDAR sweeps, 28 annotation classes, and 37 semantic segmentation labels. **Udacity** Self Driving Car Dataset has open-sourced access to a variety of projects for autonomous driving, including neural networks trained to predict steering angles of the car, camera mounts, and dozens of hours of real driving data. **Waymo Open** Dataset is an open-source multimodal sensor dataset, covers a wide variety of driving scenarios and environments. It contains 1K types of different segments where each segment captures 20 seconds of continuous driving, corresponding to 200K frames at 10 Hz per sensor. A summary of datasets considered in a nuScenes paper is below.

Dataset	Year	Sce-nes	Size (hr)	RGB imgs	PCs lidar††	PCs radar	Ann. frames	3D boxes	Night / Rain	Map layers	Clas-ses	Locations
CamVid [8]	2008	4	0.4	18k	0	0	700	0	No/No	0	32	Cambridge
Cityscapes [19]	2016	n/a	-	25k	0	0	25k	0	No/No	0	30	50 cities
Vistas [33]	2017	n/a	-	25k	0	0	25k	0	Yes/Yes	0	152	Global
BDD100K [85]	2017	100k	1k	100M	0	0	100k	0	Yes/Yes	0	10	NY, SF
ApolloScape [41]	2018	-	100	144k	0**	0	144k	70k	Yes/No	0	8-35	4x China
D^2-City [11]	2019	1k†	-	700k†	0	0	700k†	0	No/Yes	0	12	5x China
KITTI [32]	2012	22	1.5	15k	15k	0	15k	200k	No/No	0	8	Karlsruhe
AS lidar [54]	2018	-	2	0	20k	0	20k	475k	-/-	0	6	China
KAIST [17]	2018	-	-	8.9k	8.9k	0	8.9k	0	**Yes**/No	0	3	Seoul
H3D [61]	2019	160	0.77	83k	27k	0	27k	1.1M	No/No	0	8	SF
nuScenes	2019	**1k**	5.5	**1.4M**	**400k**	**1.3M**	**40k**	1.4M	**Yes/Yes**	**11**	**23**	Boston, SG
Argoverse [10]	2019	113†	0.6†	490k†	44k	0	22k†	993k†	**Yes/Yes**	2	15	Miami, PT
Lyft L5 [45]	2019	366	2.5	323k	46k	0	**46k**	1.3M	No/No	7	9	Palo Alto
Waymo Open [76]	2019	**1k**	5.5	1M	200k	0	**200k‡**	**12M‡**	**Yes/Yes**	0	4	3x USA
A*3D [62]	2019	n/a	**55**	39k	39k	0	**39k**	230k	**Yes/Yes**	0	7	SG
A2D2 [34]	2019	n/a	-	-	-	0	12k	-	-/-	0	14	3x Germany

Table 1. AV dataset comparison. The top part of the table indicates datasets without range data. The middle and lower parts indicate datasets (not publications) with range data released until and after the initial release of this dataset. We use bold highlights to indicate the best entries in every column among the datasets with range data. Only datasets which provide annotations for at least *car*, *pedestrian* and *bicycle* are included in this comparison. (†) We report numbers only for scenes annotated with cuboids. (‡) The current Waymo Open dataset size is comparable to nuScenes, but at a 5x higher annotation frequency. (††) Lidar pointcloud count collected from *each lidar*. (**) [41] provides static depth maps. (-) indicates that no information is provided. SG: Singapore, NY: New York, SF: San Francisco, PT: Pittsburgh, AS: ApolloScape.

Table 10 - Dataset summaries from the NuScences paper, processing of standard datasets can be found in MMdetection3D. The paper assesses methods applied to popular available datasets [22].

Open dataset candidates for testing within ROADVIEW DRL (MoSoCow) must / should / could fulfil.

1. **Must**
   - Open licensing terms
   - Include LiDAR, Camera and IMU measurements
2. **Should**
   - Include RADAR, LiDARs, VISIBLE Camera(s), Thermal Cameras, IMU
   - Have a toolkit to read, incorporate or indicate how to use the data
   - Have some data in adverse weather
3. **Could**
   - Gather data from trucks rather than cars
   - Appear in places such as paperswithcode.com

3 external candidates have been selected for testing within ROADVIEW as extern sources. The sources are available for non-commercial, and research purposed. Table 11 shows the datasets we selected for further investigation. Basically, a subset from the longer nuScenes paper state of the art survey.

Priority	Dataset	License	SDKs	Dataset Download
First	nuScenes	NuScenes Dataset Agreement, primarily for non-commercial academic use.	https://colab.research.google.com/github/nutonomy/nuscenes-devkit/	https://www.nuscenes.org/nuscenes#data-collection
Second	Kitty	Creative Commons Attribution-Non-Commercial-Share Alike 3.0 License	https://medium.com/multisensory-data-training/import-and-export-your-3d-point-cloud-data-in-kitti-format-with-xtreme1-sdk-toolkit-4e74c3ce3b1c	https://www.cvlibs.net/datasets/kitti/raw_data.php
Third, used in mmDetection 3D + corruptions	Waymo Open Dataset	Waymo Open Dataset License Agreement.	https://github.com/waymo-research/waymo-open-dataset	https://waymo.com/open/download/ (Redirects to Google)
(Optional)	Oxford Radar RobotCar Dataset	Creative Commons Attribution-Non-Commercial-Share Alike 4.0 International License (CC BY-NC-SA 4.0).	https://github.com/ori-mrg/robotcar-dataset-sdk	https://oxford-robotics-institute.github.io/radar-robotcar-dataset/downloads

Table 11 - Open-source datasets for use in autonomous vehicles and our processing priority

A NuScenes dataset example (Colab link)

An example from the nuScenes dataset + SDK is show below. A Notebook and execution environment available at the link in the heading. A mini dataset is available for experimentation, link. The miniset example consists of 23 categories, 8 attributes, 4 visibilities, 12 sensors, 31206 ego poses, (movement of the measurement vehicle itself), 8 logs, 10 scenes, 404 samples, 31206 sample data, 18538 sample annotations and 4 maps. Some simple examples, scene-0061, ‘Parked truck’, construction, intersection, turn left, following a van’, see Figure 17 below.

CAM_FRONT

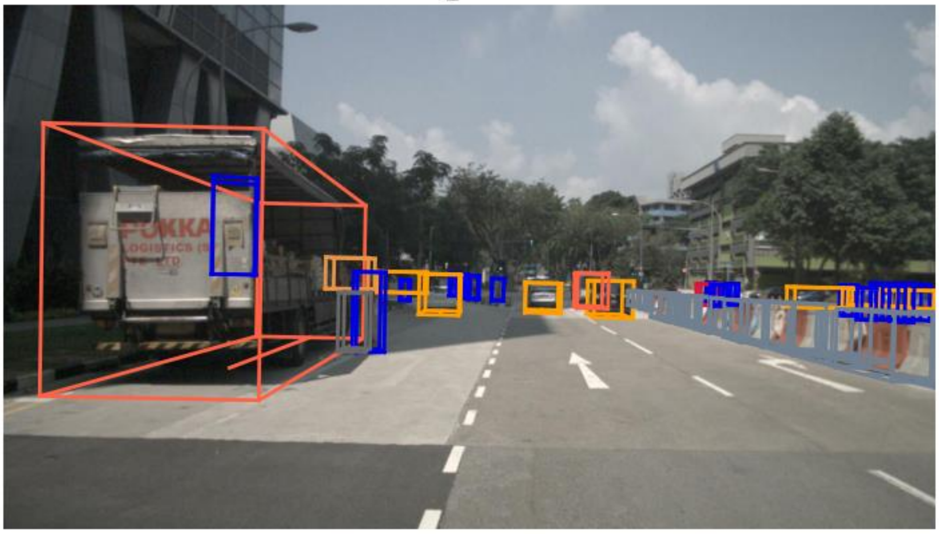

Figure 17 - nuScenes example I (https://github.com/nutonomy/nuscenes-devkit)

```
my_annotation_token = my_sample['anns'][18]
my_annotation_metadata =  nusc.get('sample_annotation', my_annotation_token)
my_annotation_metadata
```

Produces:

{'token': '83d881a6b3d94ef3a3bc3b585cc514f8',
 'sample token': 'ca9a282c9e77460f8360f564131a8af5',
 'instance token': 'e91afa15647c4c4994f19aeb302c7179',
 'visibility token': '4',
 'attribute tokens': ['58aa28b1c2a54dc88e169808c07331e3'],
 'translation': [409.989, 1164.099, 1.623],
 'size': [2.877, 10.201, 3.595],
 'rotation': [-0.5828819500503033, 0.0, 0.0, 0.812556848660791],
 'prev': '',
 'next': 'f3721bdfd7ee4fd2a4f94874286df471',
 'num_LiDAR_pts': 495,
 'num_radar_pts': 13,
 'category name': 'vehicle.truck'}

Or as a rendered image, Figure 18 below shows the same camera in front, a bounding box and the category name (in the key-value above).

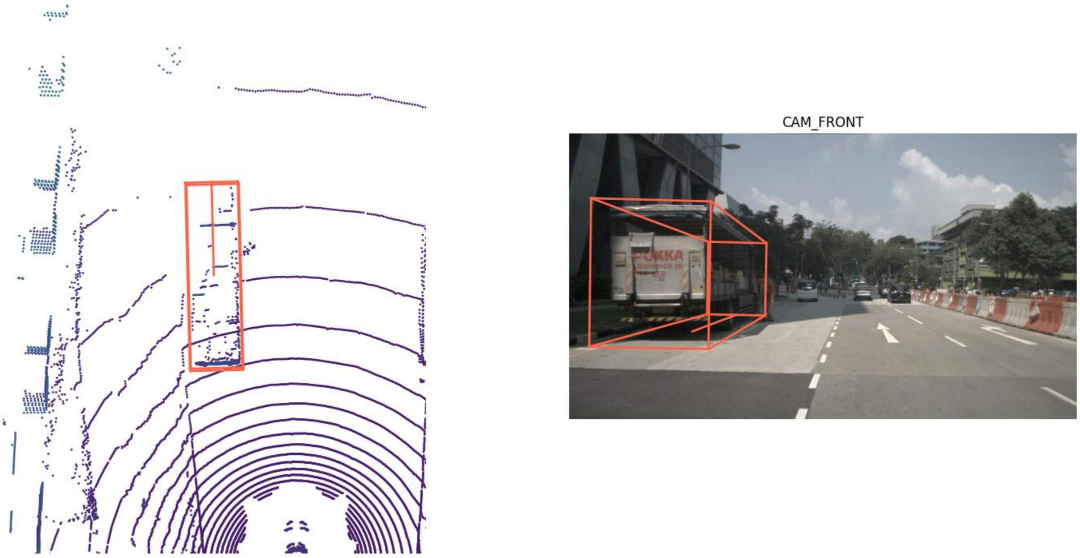

Figure 18 - nuScenes example II (https://github.com/nutonomy/nuscenes-devkit)

Below in Figure 19 are the functions of all corruptions in 3D object detection. The 3D corruptions project is built upon MMDetection3D and OpenPCDet with code modifications. The authors identify 32 LiDAR corruptions and 14 camera corruptions. They test the impact of the corruptions on the Kitty dataset [25].

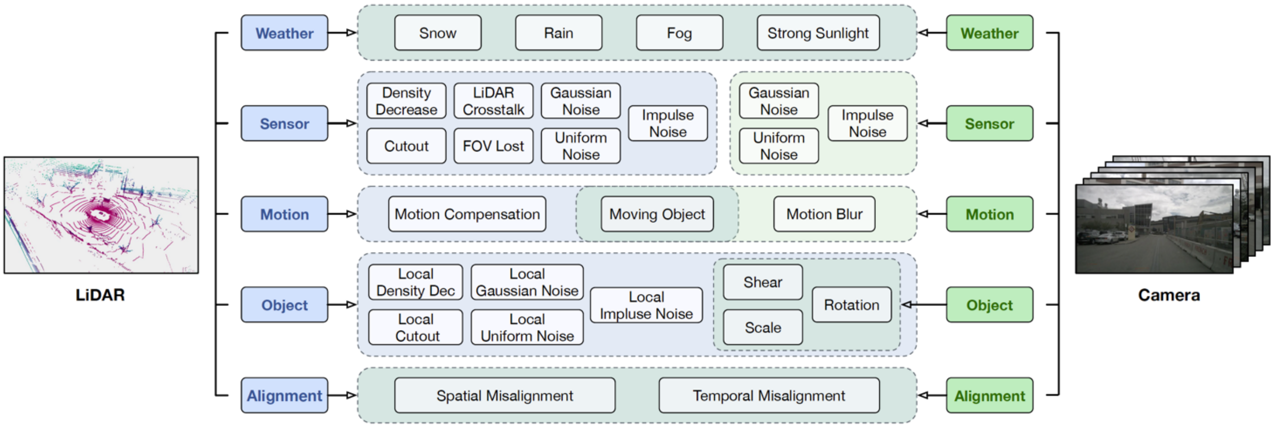

Figure 19 - Benchmarking Robustness of 3D Object Detection to Common Corruptions in Autonomous Driving

MMdetection3d

An interesting library is mmdetection3d. Not as cited as the others, but with over 1500 forks it is a feature rich dataset for experimentation, in paperswithcode, there are 63 such papers. MMDetection3D is an open-source object detection toolbox based on PyTorch, towards the next-generation platform for general 3D detection. It is a part of the OpenMMLab project. It supports multi-modality/single-modality detectors out of box using detectors including MVXNet, VoteNet, PointPillars, etc. It directly supports popular indoor and outdoor 3D detection datasets, including ScanNet, SUNRGB-D, Waymo, nuScenes, Lyft, and KITTI. For the nuScenes dataset, they also support the nuImages dataset. All the 300+ models and methods of 40+ papers as well as modules supported in MMDetection3D can be trained or used. The many authors claim it trains faster than other codebases. Like MMDetection3D and MMCV, MMDetection3D can also be used as a library to support different projects. Note it can detect images using pixel-based classic bounding boxes, semantic segmentation as well as panoptic image recognition.

Computer-generated images (simulating weather)

An alternative to using lots of data, plus expense of buying and instrumenting sensors as well as measuring, storing, processing, and driving around areas, one can generate the scenes as a sensor would see. As examples of visible images that can be generated, see Figure 20 below for examples.

Snow: Light, heavy

Snowfall: Light, Medium, heavy

Time of day: Winter daytime, morning, night

Figure 20 - Machine-generated scenes from WP2

Three ways forward using these images, is to

1. Calibration of data quality
    a. Generate a perfect image (DRL score 9)
    b. Generate a very noisy, blurred image (DRL score 1)
2. Backtrace from object misclassification to pixels
3. Generate scenes that we do not have data for.

This is work in progress and future deliverables will contain some of the results of using artificial or simulated weather conditions.

Part III: Computing & resources

This part of D4.5 is concerned with the computing aspects of the datasets and their quality. This includes the storage format, de-facto standards and alternatives, a little about the costs.

The S3 system

S3, or Amazon Simple Storage Service, is a scalable cloud storage service provided by Amazon Web Services (AWS). It does not have a specific "S3 data format." Rather, it allows you to store and retrieve various types of files, objects, or data in a highly available and durable manner. The files and objects stored in Amazon S3 can be in any format, such as text, images, videos, or binary data. We use S3 to store datasets for collaboration.

Amazon S3 organises data in a hierarchical structure with the following components:

- Buckets: These are top-level containers that store your data objects. Each bucket has a unique name within the S3 service, and you can create as many buckets as needed.
- Objects: Objects are the individual data files or items stored in S3 buckets. Each object consists of a key (a unique identifier within a bucket), data (the actual content of the object), and metadata (additional information about the object).

Amazon S3 supports various features such as versioning, lifecycle policies, access control, and data transfer acceleration. You can interact with Amazon S3 through the AWS Management Console, AWS CLI, SDKs, or RESTful APIs. They are suitable for ROADVIEW due to speed, relatively simplicity and the ability to add Access Control Lists (ACLs) to Buckets.

ROS bags

ROS bags are a file format used in the Robot Operating System (ROS), a flexible and open-source framework for developing robotics software. ROS bags provide a way to record, store, and play back data generated by ROS nodes during the operation of a robot or a simulation. This data can include sensor readings, messages exchanged between nodes, and other types of information generated within the ROS environment.

ROS bags use the .bag file extension and are designed to be an efficient and flexible format for storing large amounts of time-stamped data. They are especially useful for:

- Logging data: ROS bags can be used to record sensor data and other information during robot operation, which can be later analysed, visualised, or processed for various purposes such as debugging, testing, or performance evaluation.
- Simulation and testing: Recorded data in ROS bags can be used to replay specific scenarios or situations, enabling developers to test algorithms or tune parameters in a controlled and repeatable environment.
- Sharing data: ROS bags provide a standardised format for exchanging data between different researchers, institutions, or projects. This makes it easier to collaborate, reproduce experiments, and compare the performance of different approaches.

To work with ROS bags, you can use the rosbag command-line tool that comes with ROS. Some of the common functions in CLI format are below in `Consolas` format.

- `rosbag` **`record`**: Records data from specified topics into a new .bag file.
- `rosbag` **`play`**: Plays back the data stored in a .bag file, effectively reproducing the recorded messages and their associated timestamps.
- `rosbag` **`info`**: Displays metadata and summary information about a .bag file.
- `rosbag` **`filter`**: Extracts specific messages from a .bag file based on user-defined criteria.

In addition to the command-line tool, there are various tools and libraries available for visualising and processing data from ROS bags, such as RViz (a 3D visualisation tool for ROS) and the Python API provided by the rosbag package. The Warwick ROADVIEW dataset mentioned above is in ROS format.

Alternative storage formats

HDF5, short for Hierarchical Data Format version 5, is a versatile tool for storing and managing large numerical datasets. It organises data hierarchically, much like a file system, and supports a variety of data types. HDF5 is designed to efficiently handle very large datasets, allowing users to manipulate and process chunks of data without needing to load the entire dataset into memory. It also supports comprehensive metadata, making the data self-describing. The HDF5 file format is portable and performs well across different operating systems. Its applicability ranges across several scientific fields such as physics, astronomy, chemistry, and bioinformatics due to its ability to handle complex data structures.

Comparing the two, ROS bag is more specific to the ROS ecosystem and is optimised for storing message-passing data between software components, while HDF5 is a more general-purpose tool used in a wider range of applications for handling large and complex datasets. Conversion between HDF5 and Rosbag can be done with to and back. Alternative formats for high-speed data processing are Parquet & used by Waymo. It is a columnar binary storage file format optimised for use with big data processing frameworks like Apache Hadoop & Spark, plus many others.

The RISE ICE Cloud

One of the key features of the ICE data centre is its use of innovative cooling solutions. The facility takes advantage of the cold climate in Luleå, Sweden to cool the servers with outside air, significantly reducing the energy required for cooling. This makes the data centre more environmentally friendly and cost-effective, as it minimises energy consumption and operational costs.

The ICE data centre at RISE in Luleå serves as a testbed and research platform for various projects related to sustainable computing, including:

- Energy-efficient data storage and processing
- Renewable energy integration and smart grid solutions
- Edge computing and distributed data centres
- Cloud computing and resource optimisation
- Data centre security and resilience

By providing a state-of-the-art facility for the research and development of green data storage technologies, RISE in Luleå contributes to the global efforts towards more sustainable and energy-efficient IT infrastructure.

RISE ICE has upgraded our Ceph storage cluster with NVMe and have moved Ceph RBD based persistent storage volumes from HDD to NVMe which makes it > 50x faster. CephFS (rook-ceph-fs) volumes is still on HDD because there is much more data there. Please do some clean-up of old data and we will move that as well. RISE offers several storage solutions, one is S3-Ceph buckets, which was discussed above.

How fast is the RISE data centre?

In terms of data upload an initial speed test upload to the ice data centre showed that 650 megs (about a data CDs worth) was uploaded in about 9 secs. A DVD (4.3 GB worth of data) took about 140 seconds. The raw dataset from FGI is about 42 Gig, after processing, 60 Gig, and a merged 30 second video 2 Gig. No RADAR data is included in the FGI dataset. The S3 data format, developed by Amazon, is used in the centre, albeit an open-source implementation thereof called Ceph. After some testing, Stockholm to ice.ri.se (S3 buckets) 650 MB in 9 secs. Added NVMRe memory, this sped up access.

Costs, storage, and formats

The idea of the calculation in this table is to estimate the number of Terabytes available for the costs allocated.

UNIT	Calculation	Result	Notes
TB	0.281 SEK / hour		S3 costs. @ 2023 cost
1 year	1860 hours		Multiplying just
Project time left	3.5 years * 1860	6510 hours	6 months elapsed (4-year project)
Data costs in SEK	6510 hours * 0.281	1829 SEK	Cost per TB over project
Convert to Euro	1829	166 Euro	Cost
Project allocation	15k Euro	**90 TB (max)**	Dividing 15000 / 166

Table 12 - Storage units, costs, and formats: SEK to Euro conversion as of April 2023, (~11.7 crowns to 1 Euro). More at ice.ri.se (under-menu "pricing"). A total of 15k Euro is available to ROADVIEW, September 2022 to August 2026. Note, no computation is included in this cost, it is 'solely' storage.

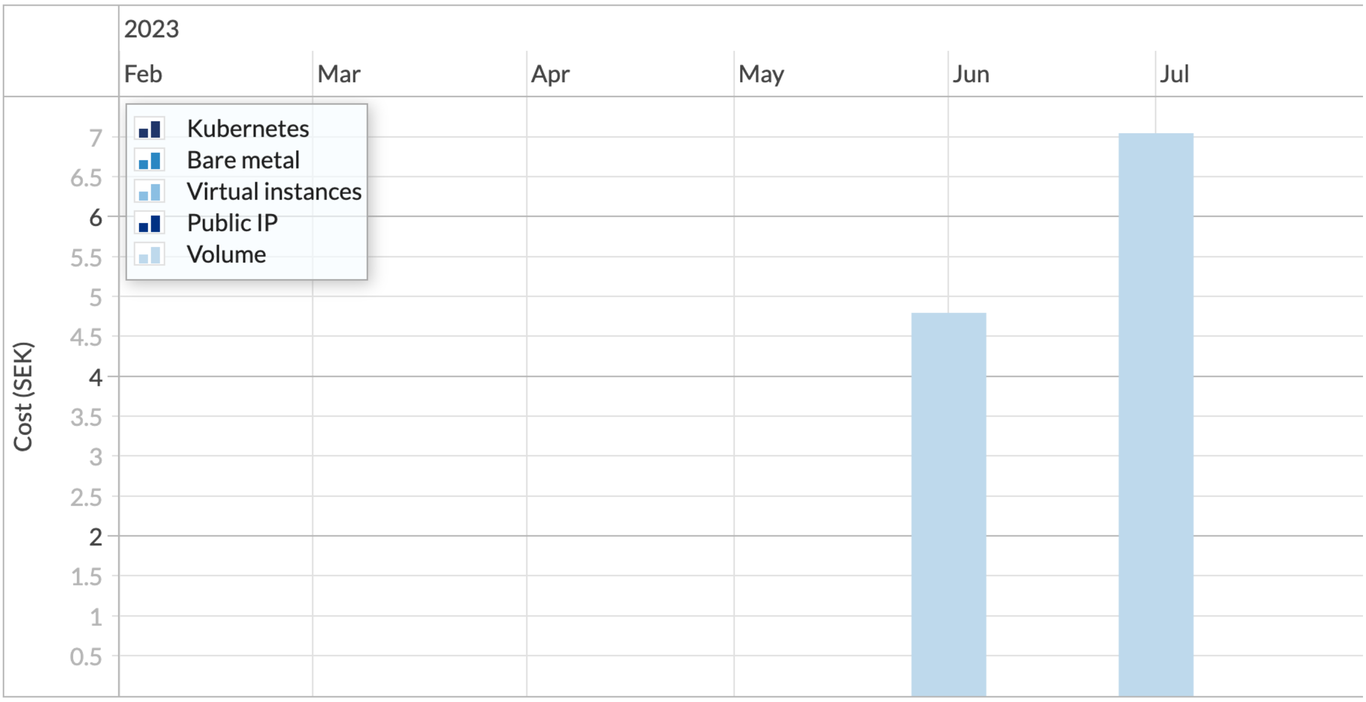

Figure 21 - Example costs for resources used in ROADVIEW (1 SEK = 0.09 Euro or 7 SEK = 0.63 EURO)

Resource	QTY	Price	Total
Kubernetes CPU	0.00 CPU / h	0.11	0.00
Kubernetes Memory	0.00 GiB / h	0.02	0.00
S3	25.10 TiB / h	0.28	7.05

Part IV: Concluding remarks

Discussion

We didn't have RADAR in the FGI dataset, but in Warwick / Carissa it exists is. It is not clear how to assess RADAR data in terms of quality now, except the positive or negative presence of an object. This is at the output stage, but where RADAR data can / could be improved we will need expert consultations. One point is how the RADAR unit is mounted and calibration. Public datasets are curated, making the datasets cleaner. This is important when releasing a dataset for experimentation. Remember 80% of the time one performs less useful work, hence the work in a public dataset, however, these may not include lots of extra work needed in a real scenario. Curated (nuScenes) versus non-curated (FGI) artificial data (Warwick simulator), often clip long range sensing (LiDAR and RADAR). Now we fuse data as separate streams, there is a philosophy where data can be collected from all sensors and fused in a single pass. This will be investigated later in the project and is one branch in the MMdetection3D dataset. Datasets for cars versus trucks. One concern from the data gathered is that we are using data gathered from vehicles, thus far. However, other WPs will produce data from a truck setup. This will enable ROADVIEW to assess the view from the correct point of view. It will also be an opportunity to create a unique dataset.

Future Work

Some of the next steps were covered in the executive summary at the beginning of this deliverable. That said, this is the initial dataset evaluation, with 1 internal (FGI) and 1 external (NuScences) dataset looked at. Most of the work has been with the image quality, and the LiDAR data from FGI using OpenPCDet. The future is to consider additional methods and datasets, from those we have selected. Development of the DRL concept will continue in the vein of additional parameters, and perhaps combining them into a class (e.g., a video). Finding scenes in public datasets that are the same or similar to those identified in WP2 will make interesting comparisons. Basically, the comparing use cases that the internal (ROADVIEW) project, with those already considered in external examples (Note: I am deliberating avoiding the EU-friendly expression "Use Case").

Conclusions

This document fulfils two purposes the readiness level of the data used in ROADVIEW. It is primarily around the sensors autonomous vehicles will carry and use, their purpose, rates, format access and quality are discussed, mostly for the first less experienced user. Subsequent deliverables will go into more detail and start to annotate data sources and ultimately assign a DRL level. This is not the final result, an indication of where to improve the data where needed. Basically, more introspection will be needed, but will provide a convenient 'score'.

ROADVIEW goes beyond the State of the Art by not only applying the DRL concept (moving from Lawrence's three bands to DRL) to the project datasets but also by using the largest set of 'tests' for each dataset. Data input and output rates will be managed by machine learning pipelines and complex issues, such as batching streamed data. ROADVIEW follows a holistic approach from the sensor hardware to the perception models presented to the decision-making system. The context of each dataset as well as their usage in the ROADVIEW system integration will define all data quality assessments.

Curated datasets such as nuScenes provide high quality data, however, are to some degree curated, for example the distance of the LiDARs is reduced to produce clear, less noisy point clouds. The real world and data are more complex with longer distances being detected.

References

[1] Lawrence, N. D. (2017). Data Readiness Levels. Available from link.

[2] Selvans, Z. (2021, June 17). Automated Data Wrangling. Catalyst Cooperative. https://catalyst.coop/2021/05/23/automated-data-wrangling.

[3] Rekatsinas, T., et al. (2017). HoloClean: Holistic Data Repairs with Probabilistic Inference. arxiv.org/pdf/1702.00820.pdf, https:/github.com/HoloClean/holoclean.

[4] Aksoy, E. E., Baci, S., & Cavdar, S. (2020, October). Salsanet: Fast Road and vehicle segmentation in LiDAR point clouds for autonomous driving. In 2020 IEEE Intelligent Vehicles Symposium (IV) (pp. 926-932). IEEE.

[5] Cortinhal, T., Tzelepis, G., & Aksoy, E. E. (2020). SalsaNext: fast, uncertainty-aware semantic segmentation of LiDAR point clouds for autonomous driving. arXiv preprint arXiv:2003.03653.

[6] Cortinhal, T., Kurnaz, F., & Aksoy, E. E. (2021). Semantics-aware multi-modal domain translation: From LiDAR point clouds to panoramic color images. In Proceedings of the IEEE/CVF International Conference on Computer Vision (pp. 3032-3048).

[7] Maanpää, J., Taher, J., Manninen, P., Pakola, L., Melekhov, I., & Hyyppä, J. (2021, January). Multimodal End-to-End Learning for Autonomous Steering in Adverse Road and Weather Conditions. In 2020 25th International Conference on Pattern Recognition (ICPR) (pp. 699-706). IEEE.

[8] Bijelic, M., Gruber, T., & Ritter, W. (2018, June). A benchmark for LiDAR sensors in fog: Is detection breaking down? In 2018 IEEE Intelligent Vehicles Symposium (IV) (pp. 760-767). IEEE.

[9] Pfeuffer, A., & Dietmayer, K. (2019, July). Robust semantic segmentation in adverse weather conditions by means of sensor data fusion. In 2019 22nd International Conference on Information Fusion (FUSION) (pp. 1-8). IEEE.

[10] Pfeuffer, A., & Dietmayer, K. (2018, July). Optimal sensor data fusion architecture for object detection in adverse weather conditions. In 2018 21st International Conference on Information Fusion (FUSION) (pp. 1-8). IEEE.

[11] Bijelic, M., Gruber, T., Mannan, F., Kraus, F., Ritter, W., Dietmayer, K., & Heide, F. (2020). Seeing through fog without seeing fog: Deep multimodal sensor fusion in unseen adverse weather. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 11682-11692).

[12] Jianqing Wu; Hao Xu; Jianying Zheng; Junxuan Zhao. Automatic Vehicle Detection with Roadside LiDAR Data Under Rainy and Snowy Conditions, IEEE Intelligent Transportation Systems Magazine, Volume: 13, Issue: 1, Spring 2021, link.

[13] Teja Vattem, George Sebastian; Luka Lukic, Rethinking LiDAR Object Detection in adverse weather conditions, 2022 International Conference on Robotics and Automation (ICRA), link.

[14] Alban Marie, Karol Desnos, Luce Morin and Lu Zhang, Metrics for Semantic Segmentation in a Machine-to-Machine Communication Scenario, 15th International Conference on Quality of Multimedia Experience (QoMEX) Evaluation of Image Quality Assessment, QoMEX, 2023.

[15] Richard Marcus, Niklas Knoop, Bernhard Egger, Marc Stamminger, A Lightweight Machine Learning Pipeline for LiDAR-simulation, 2022, arXiv.

[16] Jiyang Chen, Simon Yu, Rohan Tabish, Ayoosh Bansal, Shengzhong Liu, Tarek Abdelzaher, and Lui Sha, LiDAR Cluster First and Camera Inference Later: A New Perspective Towards Autonomous Driving, arXiv.

[17] Richard Marcus, Niklas Knoop, Bernhard Egger, Marc Stamminger, A Lightweight Machine Learning Pipeline for LiDAR-simulation, arXiv.

[18] Claudio Michaelis, Benjamin Mitzkus, Robert Geirhos, Evgenia Rusak, Oliver Bringmann, Alexander S., Ecker Matthias Bethge, Wieland Brendel, Benchmarking Robustness in Object Detection: Autonomous Driving when Winter is Coming, arXiv.

[19] LiDAR Principles, Processing and Applications in Forest Ecology, 2023.

[20] Benchmarking Robustness of 3D Object Detection to Common Corruptions in Autonomous Driving https://github.com/thu-ml/3D_Corruptions_AD.

[21] Simple, self-contained open-source project for LiDAR-based 3D object detection. https://github.com/open-mmlab/OpenPCDet.

[22] Caesar, Holger, et al. "nuscenes: A multimodal dataset for autonomous driving." *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2020.

[23] E. Yurtsever, J. Lambert, A. Carballo, and K. Takeda, "A Survey of Autonomous Driving: Common Practices and Emerging Technologies," in *IEEE Access*, vol. 8, pp. 58443-58469, 2020, doi: 10.1109/ACCESS.2020.2983149.

[24] Torres Vega, Maria & Sguazzo, Vittorio & Mocanu, Decebal & Liotta, Antonio. (2016). *An experimental Survey of No-Reference Video Quality Assessment Methods*. International Journal of Pervasive Computing and Communications. 12. 10.1108/IJPCC-01-2016-0008.

[25], Yinpeng Dong, Caixin Kang, Jinlai Zhang, Zijian Zhu, Yikai Wang, Xiao Yang, Hang Su, Xingxing Wei, Jun Zhu, Benchmarking Robustness of 3D Object Detection to Common Corruptions, CVPR2023. Link.

[26] Deepti Ghadiyaram, Chao Chen, Sasi Inguva, and Anil Kokaram, A No Reference Video Quality Predictor for compression and scaling artifacts. Link.

[27] Julius Pesonen**,** Pixelwise Road Surface Slipperiness Estimation for Autonomous Driving with Weakly Supervised Learning, Aalto University School of Science, 2023.

[28] Vargas J, Alsweiss S, Toker O, Razdan R, Santos J. An Overview of Autonomous Vehicles Sensors and Their Vulnerability to Weather Conditions. Sensors. 2021; 21(16):5397. https://doi.org/10.3390/s21165397.

[29] T. Beemelmanns, Q. Zhang, and L. Eckstein. Multicorrupt: A multi-modal robustness dataset and benchmark of lidar-camera fusion for 3d object detection, 2024.